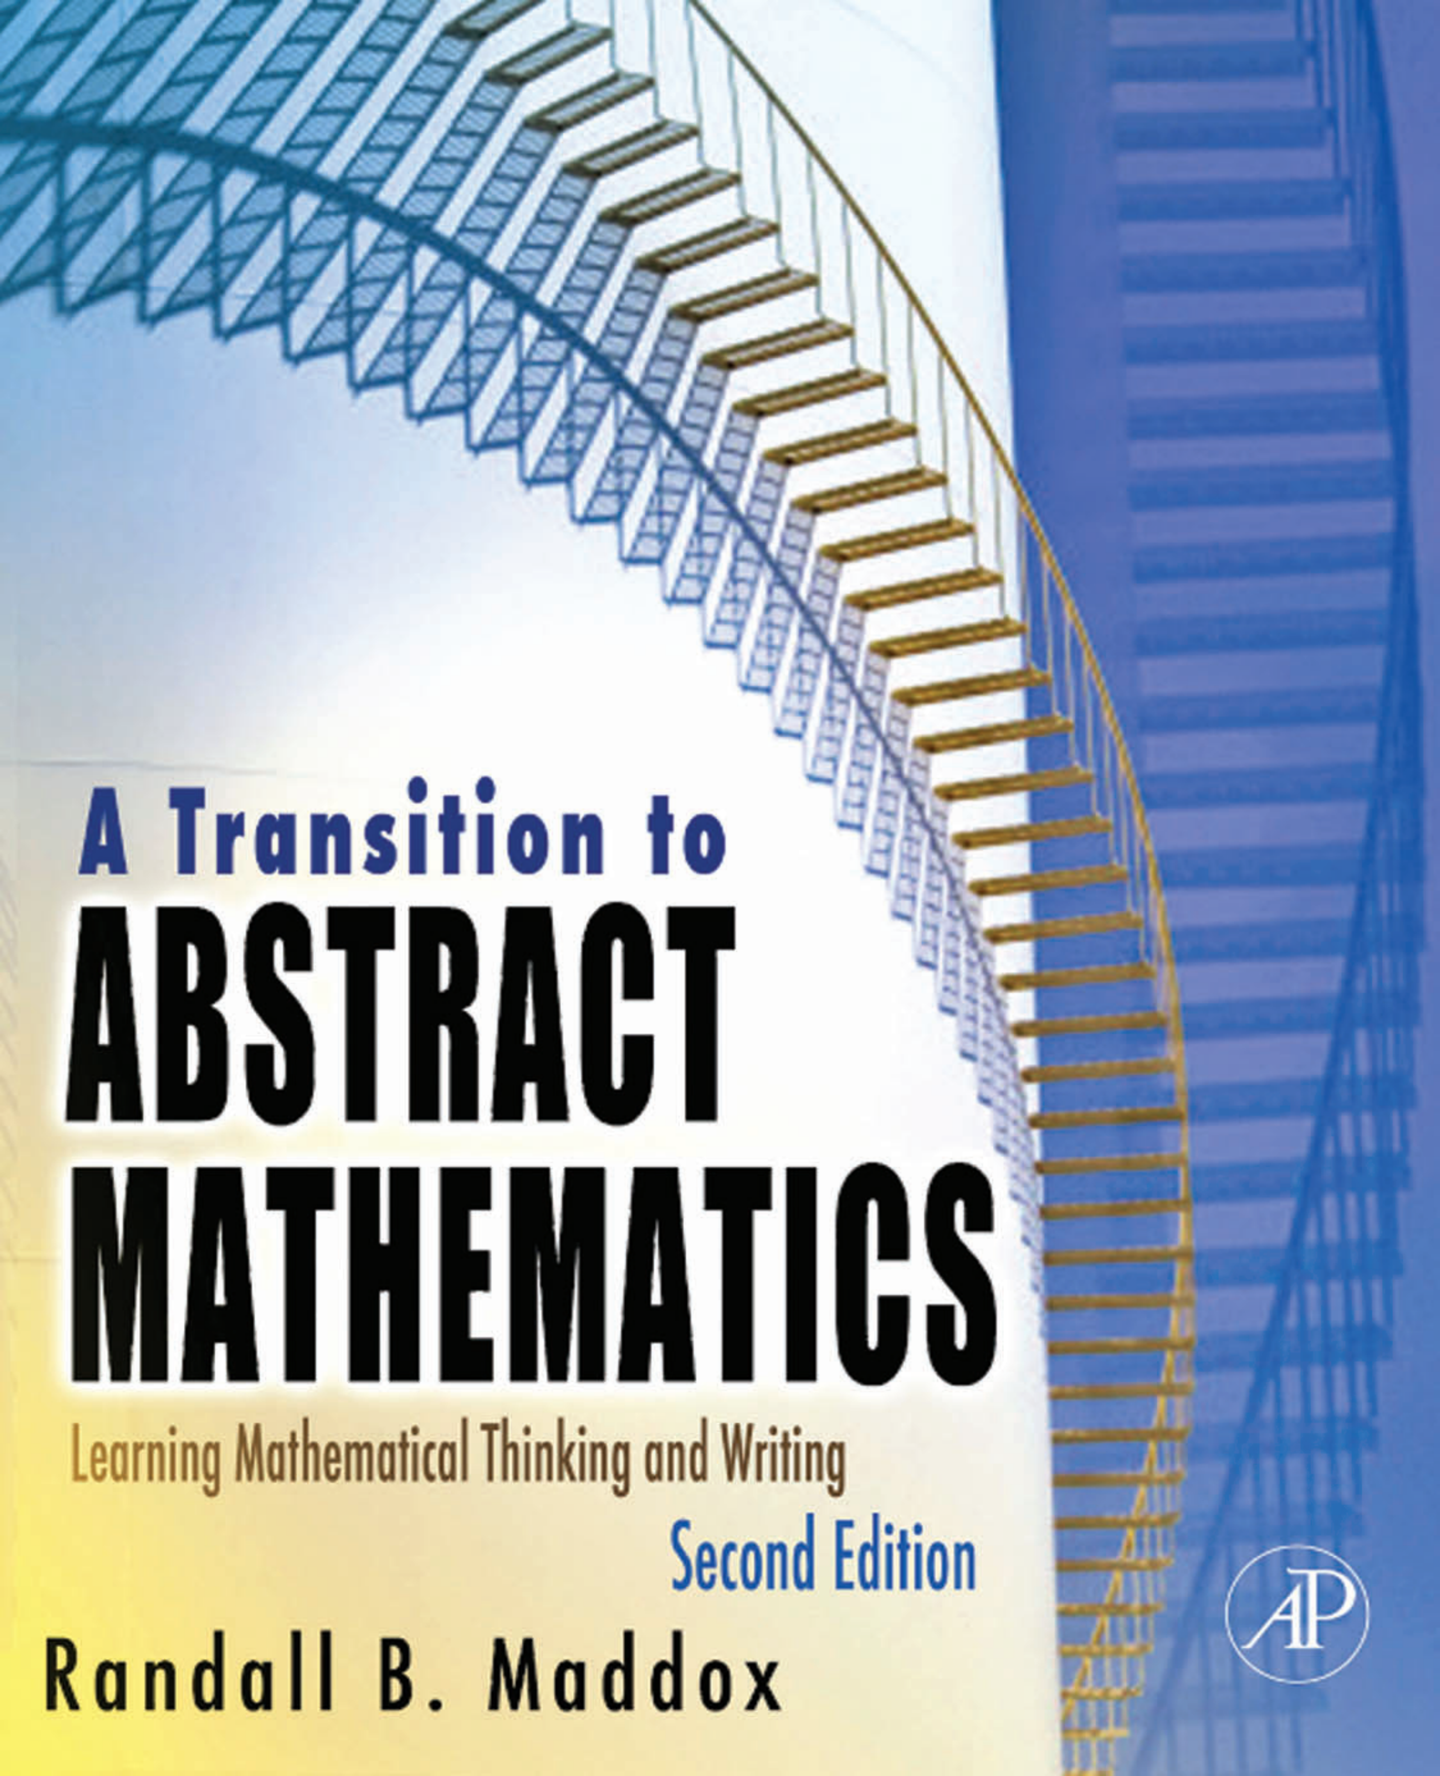

A Transition to ABSTRACT MATHEMATICS

Learning Mathematical Thinking and Writing

Second Edition

Randall B. Maddox

A Transition to Abstract Mathematics

Mathematical Thinking and Writing

Second Edition

A Transition to Abstract Mathematics

Mathematical Thinking and Writing

Second Edition

Randall B. Maddox

Pepperdine University

ELSEVIER

AMSTERDAM • BOSTON • HEIDELBERG • LONDON
NEW YORK • OXFORD • PARIS • SAN DIEGO
SAN FRANCISCO • SINGAPORE • SYDNEY • TOKYO

Academic Press is an imprint of Elsevier

Elsevier Academic Press
30 Corporate Drive, Suite 400, Burlington, MA 01803, USA
525 B Street, Suite 1900, San Diego, California 92101-4495, USA
84 Theobald's Road, London WC1X 8RR, UK

This book is printed on acid-free paper. ∞

Library of Congress Cataloging-in-Publication Data
Maddox, Randall B.
A transition to abstract mathematics: learning mathematical thinking and writing/Randall B. Maddox. – 2nd ed.
p. cm.
ISBN 978-0-12-374480-7 (hardcover: acid-free paper) 1. Proof theory. 2. Logic, Symbolic and mathematical. I. Maddox, Randall B. Mathematical thinking and writing. II. Title.
QA9.54.M34 2009
511.3′6–dc22
2008027584

British Library Cataloguing-in-Publication Data
A catalogue record for this book is available from the British Library.

ISBN 13: 978-0-12-374480-7

For all information on all Elsevier Academic Press publications visit our Web site at *www.elsevierdirect.com*

Printed in the United States of America
08 09 10 9 8 7 6 5 4 3 2 1

For Topo

my little mouse

This page intentionally left blank

Contents

This page intentionally left blank

Why Read This Book?

One of Euclid's geometry students asked a familiar question more than 2000 years ago. After learning the first theorem, he asked, "What shall I get by learning these things?" Euclid didn't have the kind of answer the student was looking for, so he did what anyone would do—he got annoyed and sarcastic. The story goes that he called his slave and said "Give him threepence since he must make gain out of what he learns."[1]

It's a familiar question: "So how am I ever gonna use this stuff?" I doubt that anyone has ever come up with a good answer, because it's really the wrong question. The first question is not what you're going to do with this stuff, but what this stuff is going to do with you.

This book is not a computer users' manual that will make you into a computer industry millionaire. It's not a collection of tax law secrets that will save you thousands of dollars in taxes. It's not even a compilation of important mathematical results for you to stack on top of the other mathematics you have learned. Instead, it's an entrance into a new kingdom, the world of mathematics, where you learn to think and write as the inhabitants do.

Mathematics is a discipline that requires a certain type of thinking and communicating that many appreciate but few develop to a great degree. Developing these skills involves dissecting the components of mathematical language, analyzing their structure, and seeing how they fit together. Once you have become comfortable with these principles, then your own style of mathematical writing can begin to shine through.

Writing mathematics requires a precision that seems a little stifling at first. It might feel like some pedant is forcing you to use a limited set of words and phrases to express the things you already see clearly with your own mind's eye. Be patient. In time you will see how adapting to the culture of mathematics and adopting its style of communicating will shape all your thinking and writing. You will see your skills of critical analysis become more developed and polished. My hope is

[1] T. L. Heath, *A History of Greek Mathematics* 1 (Oxford, 1931).

that these skills will influence the way you organize and present your thoughts in everything from English composition papers to late-night bull sessions with friends.

Here is an analogy of what the first principles of this book will do to you. Consider a beginning student of the piano. Music is one of the most creative disciplines, and our piano student has been listening to Chopin for some time. She knows she has a true ear and intuition for music. However, she must begin at the piano by playing scales over and over. These exercises develop her ability to use the piano effectively in order to express the creativity within her. Furthermore, these repetitive tasks familiarize her with the structure of music as an art form, and actually nurture and expand her capacity to express herself in original and creative ways through music. Then, once she has mastered the basic technical skills of hitting the keys, she understands more clearly how enjoyable music can be. She learns this truth: The aesthetic elements of music cannot be fully realized until the technical skills developed by rote exercises have been mastered and can be relegated to the subconscious.

Your first steps to becoming a mathematician are a lot like those for our pianist. You are going to be introduced to the building blocks of mathematical structure, and then you will practice on the precision required to communicate mathematics correctly. The drills you perform in this practice will help you see mathematics as a creative discipline and equip you to appreciate its beauty.

Think of this course as a bicycle trip through a new country. The purposes of the trip are:

- To familiarize you with the territory
- To equip you to explore it on your own
- To give you some panoramic views of the countryside
- To teach you to communicate with the inhabitants
- To help you begin to carve out your own niche

If you're willing to do the work, I promise you will enjoy the trip. Sometimes the hills are steep and the pedaling is tough. Be persistent, knowing that it's worth the effort. You will come back a different person, for this material will have done something with you. Then you'll understand that Euclid really got it right after all, and you will appreciate why his witty response is still fresh and relevant after these 2000 years.

Preface

A Transition to Abstract Mathematics was written under the assumption that students do not yet know how to read upper level mathematics texts. Since the primary purpose of the book is to teach students to write with formal rigor, and since I naturally presume they do not yet appreciate exposition written in that form, two overriding features of style defined the first edition: a loose and informal expository style of writing, and an airtight composition and organization of the logic, so that no student could ever say that any necessary detail had been overlooked or omitted. Consequently, the scope of the first edition was rather narrow and forward focused, where every exercise had an important role in the story and there were no characters too peripheral to the plot.

I believe the second edition maintains the benefits of the first edition's features but is improved in several ways. First, the exposition is still written to the student, but it is tighter and more efficient than before. Second, there are many more exercises than in the first edition. Many of these are essential in that they are the logical basis of later results. The *Instructor's Guide and Solutions Manual* points out which exercises simply must be either assigned or at least discussed because they undergird later results. Others may be assigned, discussed casually, or omitted altogether.

A third and major change to the second edition is that exercises are now integrated into the flow of the material instead of being placed at the end of each section. I believe this arrangement has several advantages. It better facilitates the students' understanding of how the mathematics is built, one step at a time, because it requires their continual participation in that process at every step. In the second edition, the text speaks clearly to the students and then presents them with exercises right on the heels of every new concept. It also should make daily course organization easier for the instructor, in that it is always clear which exercises may be assigned after a particular day's class meeting.

Other changes to the second edition include a reorganization of the material that comprised Chapter 2 in the first edition. Introductory proof-writing material on set and real number properties has now been divided into two chapters, and the order of the material basically reversed from the first edition. Thus the students' first theorems involve basic algebraic properties of numbers, which might be a simpler place for them to begin to write proofs than set properties. Chapter 1 now includes a section that enumerates different techniques of proof writing, with plenty of examples but no expectation that a student yet knows how or in what circumstances to employ these techniques. Finally, with exercises integrated into the exposition, certain sections that were quite long in the first edition have now been divided into more sections of more manageable length.

Preface to the First Edition

This text is written for a "transition course" in mathematics, where students learn to write proofs and communicate with a level of rigor necessary for success in their upper level mathematics courses. To achieve the primary goals of such a course, this text includes a study of basic principles of logic, techniques of proof, and fundamental mathematical results and ideas (sets, functions, properties of real numbers), though it goes much further. It is based on two premises: that the most important skill students can learn as they approach the cusp between lower and upper level courses is how to compose clear and accurate mathematical arguments, and that they need more help in developing this skill than they would normally receive by diving into standard upper level courses. By emphasizing how one writes mathematical prose, it is also designed to prepare students for the task of reading upper level mathematics texts. Furthermore, it is my hope that transitioning students in this way gives them a view of the mathematical landscape and its beauty, engaging them to take ownership of their pursuit of mathematics.

Why *This* Text?

I believe students learn best by doing. In many mathematics courses it is difficult to find enough time for students to discover through their own efforts the mathematics we would lead them to find. However, I believe there is no other effective way for students to learn to write proofs. This text is written to them in a format that allows them to do precisely this.

Two principles of this text are fundamental to its design as a tool whereby students learn by doing. First, it does not do too much for them. Proofs are included in this text for only two reasons. Most of them (especially at the beginning) are sample proofs that students can mimic as they write their own proofs to similar theorems. Students must read them because they will need the technique later. Other proofs are included because they are too much to expect of a student at this level. In most of these instances, however, some climactic detail is omitted and relegated to an exercise.

Second, if students are going to learn by doing, they must be presented with doable tasks. This text is designed to be a sequence of stepping stones placed just the right distance apart. Moving from one stone to the next involves writing a proof. Seeing how to step there comes from reading the exposition and calls on the experience that led the student to the current stone. At first, stones are very close together, and there is a lot of guidance. Progressing through the text, stones become increasingly farther apart, and the guidance gets less explicit.

I have written this text with a very deliberate trajectory of style. It is conversational throughout, though sophistication and succinctness of the exposition increase from chapter to chapter.

Organization

This text is organized in the following way. Chapter 0 spells out all assumptions to be used in writing proofs. These are not necessarily standard axioms of mathematics, and they are not presented in the context or language of more abstract mathematical structures. They are designed merely to be a starting point for logical development, so that students appreciate quickly that everything we call on is either stated up front as an assumption or proved from these assumptions. Students can probably read Chapter 0 on their own as the course begins, knowing that it is there primarily as a reference.

Part I begins with logic but does not focus on it. In Chapter 1, truth tables and manipulation of logical symbols are included to give students an understanding of mathematical grammar, of the underlying skeleton of mathematical prose, and of equivalent ways of communicating the same mathematical idea. Chapters 2–4 put these to use right away in proof writing, and allow the students to cut their teeth on the most basic mathematical ideas. These chapters will constitute most, or perhaps all, of the content of the course.

Parts II and III are two completely independent paths, the former into analysis, the latter into algebra. Like Antoni Gaudí's *Sagrada Familia*, the unfinished cathedral in Barcelona, Spain, where narrow spires rise from a foundation to give spectacular views, Parts II and III are purposefully designed to rest on the foundation of Part I and climb quickly into analysis or algebra. Many topics and specific results are omitted along the way, but Parts II and III rest securely on the foundation of Part I and allow students to continue to develop their skills at proof writing by climbing to a height where, I hope, they have a nice view of mathematics.

Flexibility

This text can be used in a variety of ways. It is suitable for use in different class settings, and there is much flexibility in the material one may choose to cover.

First, because this text speaks directly to the student, it can naturally be used in a setting where students are given responsibility for the momentum of the class. It is written so that students can read the material on their own first, then bring to class the fruits of their work on the exercises, and present these to the instructor and each other for discussion and critique. If class time and size limit the practicality of such a student-driven approach, then certainly other approaches are possible. To illustrate, we may consider three components of a course's activity and arrange them in several ways. The components are (1) the students' reading of the material, (2) the instructor's elaboration on the material, and (3) the students' work on the exercises, either to be presented in class or turned in. When I teach from this text, (1) is first, (3) follows on its heels, and (2) and (3) work in conjunction until a section is finished. Others might want to arrange these components in another order, for example, beginning with (2), then following with (1) and (3).

Which material an instructor would choose to cover will depend on the purpose of the course, personal taste, and how much time there is. Here are two broad options.

1. To proceed quickly into either analysis or algebra, first cover the material from Part I that lays the foundation. Almost all sections and exercises of Part I are necessary for Parts II and III. However, the *Instructor's Guide and Solutions Manual* notes precisely which sections, theorems, and exercises are necessary for each path, and which may be safely omitted without leaving any holes in the logical progression. Of course, even if a particular result is necessary later, one might decide that to omit its proof details does not deprive the students of a valuable learning experience. The instructor might choose simply to elaborate on how one would go about proving a certain theorem, then allow the students to use it as if they had proved it themselves.

2. Cover Part I in its entirety, saving specific analysis and algebra topics for later courses. This option might be most realistic for courses of two or three units where all the Part I topics are required. Even with this approach, there would likely be time to cover the beginnings of Parts II and/or III. This might be the preferred choice for those who do not want to study analysis or algebra with the degree of depth and breadth characteristic of this text.

This page intentionally left blank

Acknowledgments

It takes an entire team to write a book. It strikes me how a well-coordinated team can pull this feat off without ever meeting each other face to face, or in some cases, without even knowing who the other team members are. So it is with the second edition of *Transition to Abstract Mathematics*.

First, enthusiastic thanks go to the staff at Elsevier who coordinated and drove this project to its completion: editors Lauren Schultz and Gavin Becker, project manager Julie Ochs, Leah Ackerson in marketing, and cover designer Eric DeCicco. Next are the professors and students who provided valuable input as the revised manuscript took shape: Michael Coco at Lynchburg College, Jessica Knapp at Pima Community College, and Will Cousins at Pepperdine University.

Finally, my deepest appreciation goes to the many students who have made the adventurous transition to abstract mathematics under my direction over the past few years. I take endless delight in seeing them gain their mathematical legs as they learn to stand, walk, and then run on their own. Without exception, I can read in their demeanor and hear in their conversation that they have developed a new way of thinking and communicating, as well as a new level of confidence in their ability to play the mathematical game.

This page intentionally left blank

Notation and Assumptions

Suppose you have just opened a new jigsaw puzzle. What are the first things you do? First, you pour all the pieces out of the box. Then you sort through and turn them all face up, taking a quick look at each one to determine whether it is an inside or outside piece, and you arrange them somehow so that you will have an idea of where certain types of pieces can be found later. You don't study each piece in depth, nor do you start trying to fit any of them together. You merely lay all the pieces out on the table and briefly familiarize yourself with them. This is the point of the game where you set the stage, knowing that everything you will need later has been put in a place where you can find it when you need it.

In this introductory chapter, we lay out all the pieces we will use for our work in this course. It is essential that you read it now, in part because you need some preliminary exposure to the ideas, but mostly because you need to have spelled out precisely what you can use without proof in Part I, where this chapter will serve you as a reference. Give this chapter a casual but complete reading for now. You have probably encountered most of the ideas before. But don't try to remember it all, and certainly don't expect to understand everything either. That is not the point. Right now, we are just organizing the pieces. The two issues we address in this chapter are: (1) Set terminology and notation, and (2) Assumptions about the real numbers.

0.1 Set Terminology and Notation

Sets are perhaps the most fundamental mathematical entity. Intuitively, we think of a set as a collection of things, where the collection itself is regarded as a single entity. Sets may contain numbers, points in the xy-plane, functions, ice cream cones, steak knives, worms, even other sets. We will denote many of our sets with uppercase letters (A, B, C), or sometimes with scripted letters ($\mathcal{F}$, $\mathcal{S}$, $\mathcal{T}$). First, we need a way of stating whether a certain thing is or is not in a set.

Definition 0.1.1 If A is a set and x is an entity in A, we write $x \in A$, and say that x is an *element* of A. To write $x \notin A$ means that x is not an element of A.

How can you communicate to someone what the elements of a set are? There are several ways.

1. List them. If there are only a few elements in the set, you can easily list them all. Otherwise, you might start listing the elements and hope that the reader can take the hint and figure out the pattern. For example,

   (a) $\{1, 8, \pi, \text{Monday}\}$

   (b) $\{0, 1, 2, \dots, 40\}$

   (c) $\{\dots, -6, -4, -2, 0, 2, 4, 6, \dots\}$

2. Provide a description of the criteria used to define whether an entity is to be included. It works like this:

   (a) $\{x : x \text{ is a real number and } x > -1\}$
   This notation should be read "the set of all x such that x is a real number and x is greater than -1." The indeterminate x is just a symbol chosen to represent an arbitrary element of the set, so that any characteristics it must have can be stated in terms of that symbol.

   (b) $\{p/q : \ p \text{ and } q \text{ are integers and } q \neq 0\}$
   This is the set of all fractions, integer over integer, where it is expressly stated that the denominator cannot be zero.

   (c) $\{x : P(x)\}$
   This is a generic form for this way of describing a set. The expression $P(x)$ represents some specified property that x must have in order to be in the set.

Some of the sets we will use most are the following:

Empty set:	$\emptyset = \{\,\}$ (the set with no elements)
Natural numbers:	$\mathbb{N} = \{1, 2, 3, \dots\}$
Whole numbers:	$\mathbb{W} = \{0, 1, 2, 3, \dots\}$
Integers:	$\mathbb{Z} = \{\dots, -3, -2, -1, 0, 1, 2, 3, \dots\}$
Rational numbers:	$\mathbb{Q} = \{p/q : p, q \in \mathbb{Z}, q \neq 0\}$
Real numbers:	$\mathbb{R}$ (Explained in the next section)
Nonzero real numbers:	$\mathbb{R}^{\times}$

Given two sets A and B, it just might happen that all elements of A are also elements of B. We write this as $A \subseteq B$ and say that A is a *subset* of B. Equivalently, we may write $B \supseteq A$, and say that B is a *superset* of A. If A is a subset of B, but there are elements of B that are not in A, we say that A is a *proper subset* of B, and write this $A \subset B$.

Example 0.1.2 Using the sets of numbers defined previously, it follows that $\mathbb{N} \subset \mathbb{W}$ and $\mathbb{W} \subseteq \mathbb{Z}$. It is also true that $\emptyset \subseteq \mathbb{N}$. ■

Anytime we talk about a particular set, there is always a context within which the set is assumed to exist. For example, let A be the set of students enrolled in your math class, and let B be the set of first-year students at your college who are taking beginning French. One possible context for these sets is the set of all students at your college. It is as if we have a largest set from which all elements of all sets in the current discussion are taken. This largest set is called the *universal set* and is denoted U. We need a universal set in order to define the *complement* of a set.

Definition 0.1.3 Given a set A, the set A^C is called the *complement* of A and is defined as the set of all elements of U that are not in A. That is,

$$A^C = \{x : x \in U \text{ and } x \notin A\}$$

Example 0.1.4 If the real numbers are taken as the universal set, then

$$\mathbb{Q}^C = \{x : x \in \mathbb{R} \text{ and } x \notin \mathbb{Q}\}$$

This is the set of irrational numbers. ■

0.2 Assumptions about the Real Numbers

One big question we will face when we begin to write proofs is what we are allowed to assume and what we must justify with proof. The purpose of this section is to lay out all the assumptions we will make concerning the real numbers. In later chapters, we will restate these assumptions when we first need to apply them. We outline them here for the sake of reference.

0.2.1 Basic Algebraic Properties

The real numbers, as well as their familiar subsets $\mathbb{N}$, $\mathbb{W}$, $\mathbb{Z}$, and $\mathbb{Q}$, are assumed to be endowed with the relation of equality and the operations of addition and multiplication and to have the following properties. First, equality is assumed to behave in the following way.

(A1) **Properties of Equality:**

(a) For all $a \in \mathbb{R}$, $a = a$. (Reflexive)

(b) For all $a, b \in \mathbb{R}$, if $a = b$, then $b = a$. (Symmetric)

(c) For all $a, b, c \in \mathbb{R}$, if $a = b$ and $b = c$, then $a = c$. (Transitive)

The first property of addition we will assume concerns its predictable behavior, even when the numbers involved can be addressed by more than one name. For example, 3/8 and 6/16 are different names for the same number. We need to know that adding something to 3/8 will always produce the same result as adding it to 6/16. The following property is our way of stating this assumption.

(A2) **Addition is well defined:** For all $a, b, c, d \in \mathbb{R}$, if $a = b$ and $c = d$, then $a + c = b + d$.

A special case of property A2 yields a familiar principle that goes back to your first days of high school algebra. That is the fact that if $a = b$, then since $c = c$, we have that $a + c = b + c$.

(A3) **Closure property of addition:** For all $a, b \in \mathbb{R}$, $a + b \in \mathbb{R}$. That is, the sum of two real numbers is still a real number. This closure property also holds for $\mathbb{N}$, $\mathbb{W}$, $\mathbb{Z}$, and $\mathbb{Q}$.

(A4) **Associative property of addition:** For all a, b, $c \in \mathbb{R}$, $(a + b) + c = a + (b + c)$.

Addition is a *binary operation*, meaning it combines exactly two numbers to produce a single number result. If we have three numbers a, b, and c to add up, we must split the task into two steps of adding two numbers. Property A4 says it does not matter which two, a and b, or b and c, we add first. It motivates the more lax notation $a + b + c$.

(A5) **Commutative property of addition:** For all $a, b \in \mathbb{R}$, $a + b = b + a$.

(A6) **Existence of an additive identity:** There exists an element $0 \in \mathbb{R}$ with the property that $a + 0 = a$ for all $a \in \mathbb{R}$.

(A7) **Existence of additive inverses:** For all $a \in \mathbb{R}$, there exists some $b \in \mathbb{R}$ such that $a + b = 0$. Such an element b is called an *additive inverse* of a and is typically denoted $-a$ to show its relationship to a. We do *not* assume that only one such b exists.

Properties similar to A2–A7 hold for multiplication.

(A8) **Multiplication is well defined:** For all $a, b, c, d \in \mathbb{R}$, if $a = b$ and $c = d$, then $ac = bd$.

(A9) **Closure property of multiplication:** For all $a, b \in \mathbb{R}$, $ab \in \mathbb{R}$. The closure property of multiplication also holds for $\mathbb{N}$, $\mathbb{W}$, $\mathbb{Z}$, and $\mathbb{Q}$.

(A10) **Associative property of multiplication:** For all $a, b, c \in \mathbb{R}$, $(ab)c = a(bc)$.

(A11) **Commutative property of multiplication:** For all $a, b \in \mathbb{R}$, $ab = ba$.

(A12) **Existence of a multiplicative identity:** There exists an element $1 \in \mathbb{R}$ with the property that $a \cdot 1 = a$ for all $a \in \mathbb{R}$.

(A13) **Existence of multiplicative inverses:** For all *nonzero* $a \in \mathbb{R}$, there exists some $b \in \mathbb{R}$ such that $ab = 1$. Such an element b is called a *multiplicative inverse* of a and is typically denoted a^{-1} to show its relationship to a. As with additive inverses, we do not assume that only one such b exists. Furthermore, the assumption that a multiplicative inverse exists for all nonzero real numbers does *not* assume that zero has no multiplicative inverse. It says nothing about zero at all.

The next property describes how addition and multiplication interact.

(A14) **Distributive property of multiplication over addition:** For every $a, b, c \in \mathbb{R}$, $a(b + c) = (ab) + (ac) = ab + ac$, where the multiplication is assumed to be done before addition in the absence of parentheses.

Property A14 is important because it is the only link between the operations of addition and multiplication. Several important properties of the real numbers owe their existence to this relationship. For example, as we will see later, the fact that $a \cdot 0 = 0$ for every real number a is a direct result of the distributive property, and not something we simply assume.

From addition and multiplication we create the operations of subtraction and division, respectively. Knowing that additive and multiplicative inverses exist (except for 0^{-1}), we write

$$a - b = a + (-b)$$

$$a/b = a \cdot b^{-1}$$

One very important assumption we need concerns properties A6 and A12. For reasons you will see later, we need to assume that the additive identity is different from the multiplicative identity. That is, we need the assumption

(A15) $1 \neq 0$.

We will use these very basic properties to derive some other familiar properties of real numbers in Chapter 2.

0.2.2 Ordering Properties

One standard way of comparing two real numbers is with the *greater than* symbol $>$. Intuitively, we think of the statement $a > b$ as meaning that a is to the right of b on the number line. Though this is helpful, the comparison $a > b$ is actually a bit sticky. The nuts and bolts of $>$ are contained in the following. In A16, we make an assumption about how real numbers compare to zero by $>$, thus giving meaning to the terms *positive* and *negative*. Then in A17 and A18, we make some assumptions about how the positive real numbers behave.

(A16) **Trichotomy law:** For any $a \in \mathbb{R}$, exactly one of the following is true:

(a) $a > 0$, in which case we say a is *positive*

(b) $a = 0$

(c) $0 > a$, in which case we say a is *negative*

(A17) For all $a, b \in \mathbb{R}$, if $a > 0$ and $b > 0$, then $a + b > 0$. That is, the set of positive real numbers is closed under addition.

(A18) For all $a, b \in \mathbb{R}$, if $a > 0$ and $b > 0$, then $ab > 0$. That is, the set of positive real numbers is closed under multiplication.

Now we can use A16–A18 to give meaning to other statements comparing any pair of real numbers.

Definition 0.2.1 Given real numbers a and b, we say that $a > b$ if $a - b > 0$. The statement $a < b$ means $b > a$. The statement $a \geq b$ means that either $a > b$ or $a = b$. Similarly, $a \leq b$ means either $a < b$ or $a = b$.

The rest of the properties of real numbers are probably not as familiar as the ones above, but their roles in the theory of real numbers will be clarified in good time. As with the previous properties, we do not try to justify them. We merely accept them and use them as a basis for proofs. A very important property of the whole numbers is the following.

(A19) **Well-ordering principle:** Any non-empty subset of whole numbers has a smallest element. That is, if A is a non-empty set of whole numbers, then there is some number $a \in A$ with the property that $a \leq x$ for all $x \in A$. In particular, we assume that 1 is the smallest natural number.

The next property of the real numbers is a bit complicated but is indispensible in the theory of real numbers. Read it casually for the first time, but know that it will be very important in Part II of this text. Suppose A is a non-empty subset of the real numbers with the property that it is bounded from above. That is, suppose there is some real number M with the property that $a \leq M$ for all $a \in A$. For example, let $A = \{x : x^2 < 10\}$. Clearly, $M = 4$ is a number such that every $a \in A$ satisfies $a \leq M$. So 4 is an *upper bound* for the set A. There are other upper bounds for A, such as 10, 3.3, and 3.17. The point is that, among all upper bounds that exist for a set, there is an upper bound that is smallest, and it exists in the real numbers. This is stated in the following.

(A20) **Least upper bound property:** If A is a non-empty subset of the real numbers that is bounded from above, then there exists a least upper bound in the real numbers. That is, if there exists some $M \in \mathbb{R}$ with the property that $a \leq M$ for all $a \in A$, then there will also exist some $L \in \mathbb{R}$ with the following properties:

(L1) For every $a \in A$, we have that $a \leq L$, and

(L2) If N is any upper bound for A, it must be that $N \geq L$.

0.2.3 Other Assumptions

The real numbers are indeed a complicated set. The final real number properties we mention are not standard assumptions, and they deserve your attention at some point in your mathematical career. In this text, we assume them.

(A21) The real numbers can be equated with the set of all base 10 decimal representations. That is, every real number can be written in a form like 338.1898..., where the decimal might or might not terminate, and might or might not fall into a pattern of repetition. Furthermore, every decimal form you can construct represents a real number. Strangely, though, there might be more than one decimal representation for a certain real number. You might remember that $0.9999\ldots = 1$. The repeating 9 is the only case where more than one decimal representation is possible. We will assume this.

Our final assumption concerns the existence of roots of real numbers.

(A22) For every positive real number x and any natural number n, there exists a real number solution y to the equation $y^n = x$. Such a solution y is called an nth root of x. The common notation $\sqrt[n]{x}$ will be addressed in Section 3.9.

Notice we make no assumptions about how many such roots of x there are, or what their signs are. Nor do we assume anything about roots of zero or of negative real numbers. We will derive these from assumption A22.

One final comment about assumptions in mathematics is in order. In a rigorous development of any mathematical theory, some things must be assumed without proof; that is, they must be *axiomatic*, serving as a starting place for the mathematician's thinking. In a study of the real numbers, some of the assumptions A1–A22 are standard. Others would be considered standard assumptions only for some subsets of the real numbers, perhaps for the whole numbers. The mathematician would then very painstakingly apply assumptions made to the whole numbers in order to expand the same properties to all of the real numbers. One assumption in particular, A21, is a most presumptuous one. But let us make no apologies for this. After all, many of the foundational issues in mathematics were addressed very late historically, and this is not a course in the foundations of mathematics. It is a course to teach us how mathematics is done and to give us some enjoyment of that process. We choose assumptions here that likely coincide with your current idea of a reasonable place to start. In some cases, we will dig more deeply as we go, though some of the foundational work will come in your later courses.

PART I

Foundations of Logic and Proof Writing

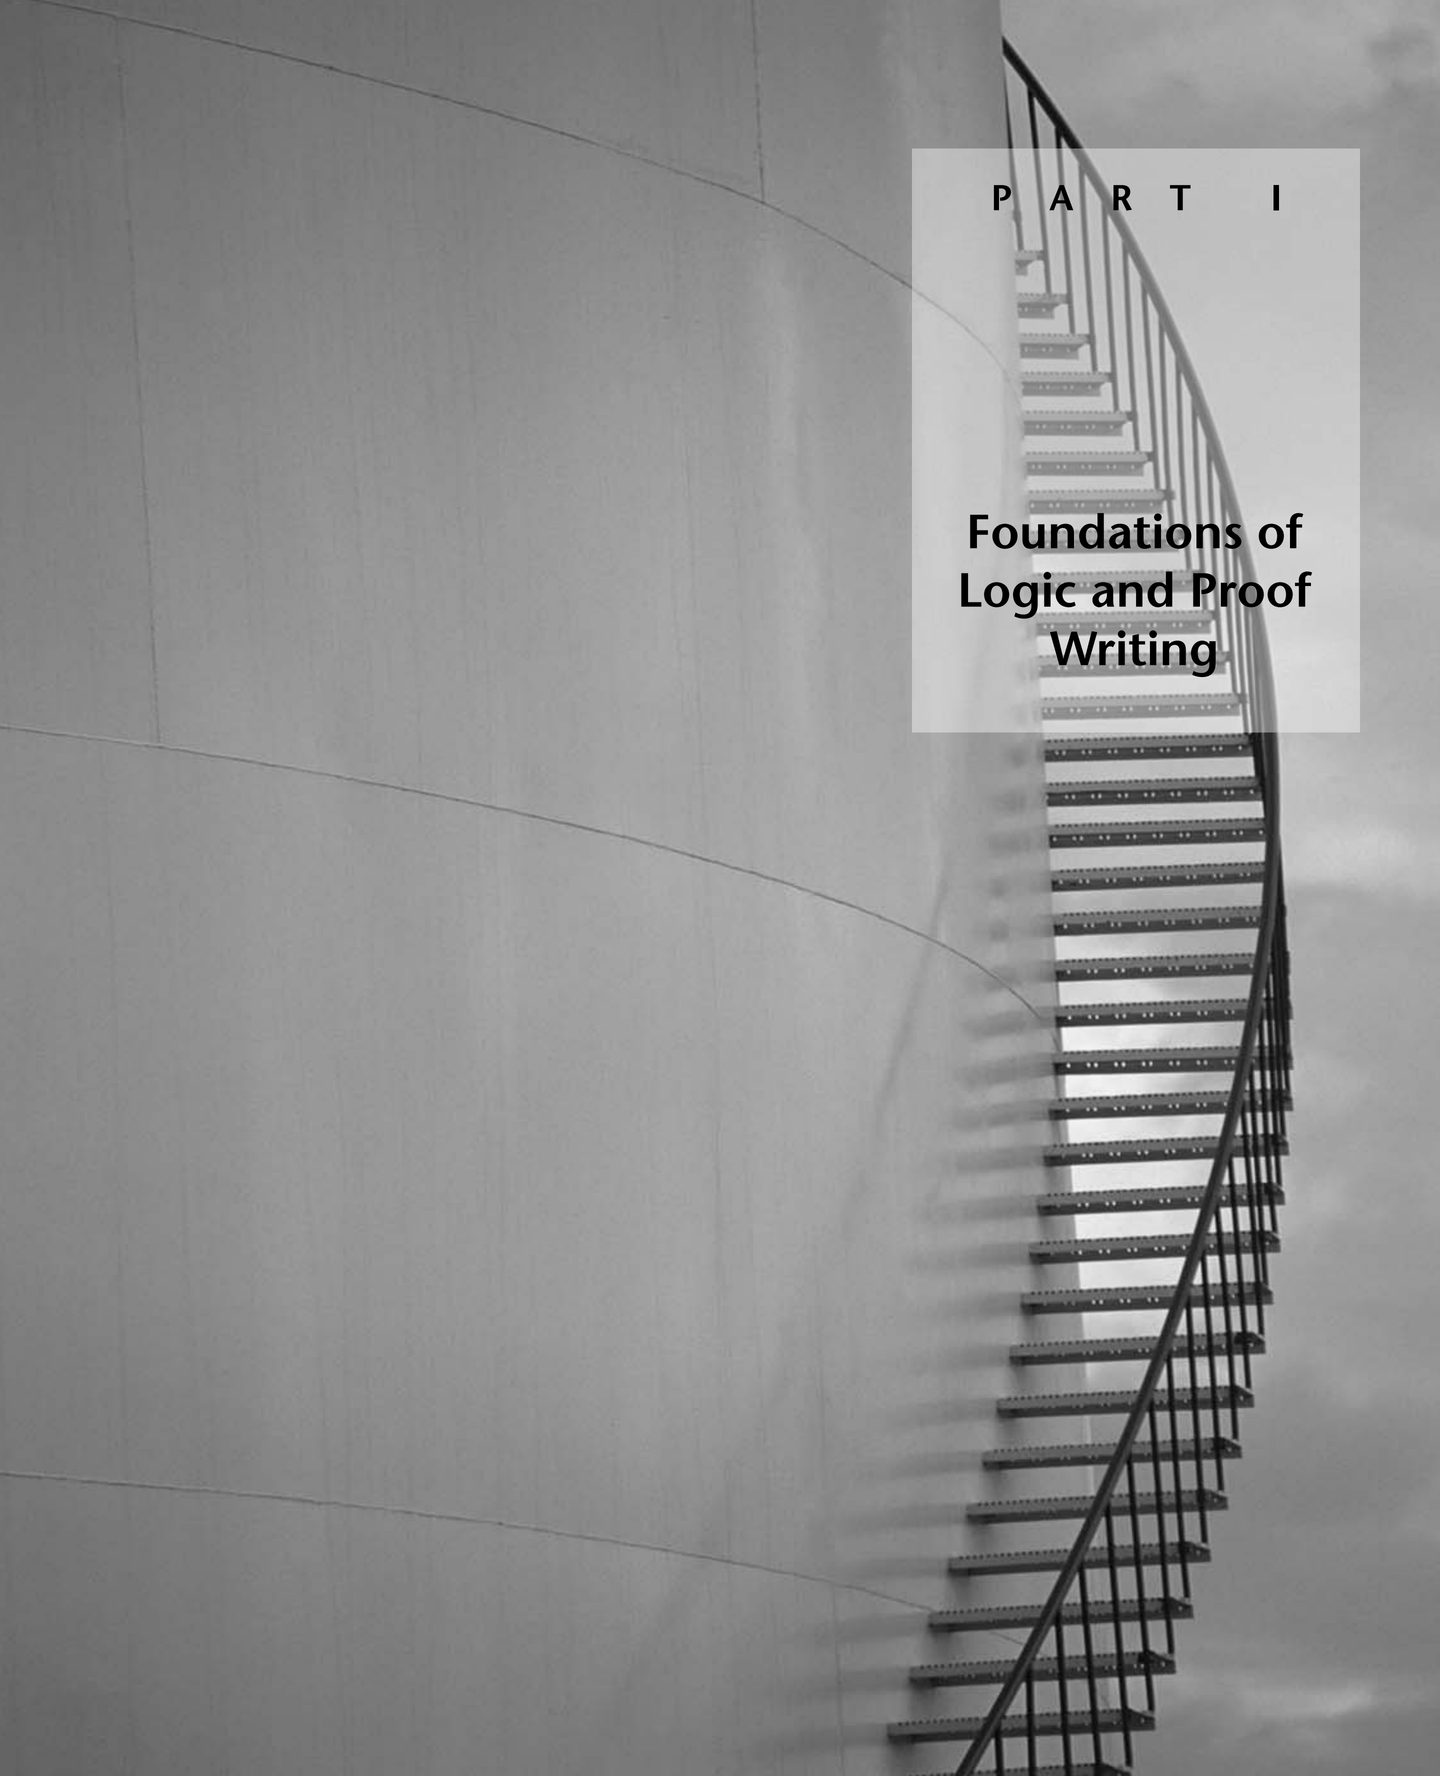

This page intentionally left blank

1

Language and Mathematics

One main purpose of this text is to develop your use of language in the context of mathematics. In this chapter, we will lay out some of the principles that govern the mathematician's very precise use of language.

1.1 Introduction to Logic

1.1.1 Statements

The first issue we address is what kinds of sentences mathematicians use as building blocks of their work. Remember from elementary school grammar, sentences are generally divided into four classes:

Declarative sentences: We also call these *statements*. Here are some examples:

1. Labor Day is the first Monday in September.
2. Earthquakes don't happen in California.
3. 3 is greater than 7.
4. 28,657 is a prime number.
5. If F_n is the nth Fibonacci number and if F_n is prime, then n is prime.

One characteristic of statements that jumps out at you is that they generally evoke a reaction like "Yeah, that's true," or "No way," or even "That could be true, but I don't know for sure." Statements 1, 4, and 5 are true, and statements 2 and 3 are false.

Imperative sentences: We would call these commands.

1. Don't wash your red bathrobe with your white underwear.
2. Knock three times on the ceiling if you want me.

Interrogative sentences: That is, questions.

1. How much is that doggy in the window?
2. Why do fools fall in love?

Exclamations:

1. What a day!
2. Lions and tigers and bears, Oh my!
3. So, like, whatever.

The building blocks of the mathematician's work are statements, but we have to be careful about exactly which declarative sentences we allow. We will define a *statement* intuitively as a sentence that can be assigned either to the class of things we would call TRUE or to the class of things we would call FALSE. Two pitfalls present themselves immediately.

First, there are paradoxes. For example, "This sentence is false" cannot be either true or false. If you think the sentence is true, then it is false. But, if it is false, then it is true. We do not want to consider paradoxes as statements.

EXERCISE 1.1.1 Suppose the barber of Seville is a man who lives in the town of Seville. Determine whether each of the following statements is a paradox.

(a) The barber of Seville shaves every man in the town of Seville who does not shave himself.

(b) The barber of Seville does not shave any man in the town of Seville who shaves himself.

(c) The barber of Seville shaves every man in Seville who does not shave himself, but he does not shave any man in Seville who does shave himself.

Second, some expressions contain what logicians call an *indeterminate*. The presence of an indeterminate in an expression excludes the expression from being a statement. For example, to say

$$x \text{ can be written as the sum of two prime numbers}$$

is not considered a statement. The use of the indeterminate x is like leaving a blank unfilled, so that the expression cannot be classified as either true or false. However, if we were to say

Every integer between 3 and 20 can be written
as the sum of two prime numbers,

then this is a statement.

Now imagine the set of all conceivable statements, and call it $\mathcal{S}$. Certainly, this set is frighteningly large and complex, but a most important characteristic of its elements is that each one can (at least in theory) be placed into exactly one of two subsets: $\mathcal{T}$ (statements called TRUE) and $\mathcal{F}$ (statements called FALSE). We want to look at relationships between these statements. Specifically, we want to pick statements from $\mathcal{S}$, change or combine them to make other statements in $\mathcal{S}$, and lay out some understandings of how the truth or falsity of the chosen statements determines the truth or falsity of the alterations and combinations. In the next part of this section, we discuss three ways of doing this:

- The negation of a statement
- A compound statement formed by joining two given statements with AND
- A compound statement formed by joining two given statements with OR

1.1.2 Negation of a Statement

We generally use p, q, r, and so forth, to represent statements symbolically. For example, define a statement p as follows:

$$p: \text{ Meghan has rented a car for today.}$$

Now consider the *negation* or *denial* of p, which we can create by a strategic placement of the word NOT somewhere in the statement. We write it this way:

$$\neg p: \text{ Meghan has not rented a car for today.}$$

If p is true, then $\neg p$ is false, and vice versa. We illustrate this in a *truth table*, Table 1.1.

p	$\neg p$
T	F
F	T

(1.1)

Definition 1.1.2 Given a statement p, we define the statement $\neg p$ (not p) to be false when p is true, and true when p is false, as illustrated in Table 1.1.

1.1.3 Combining Statements with AND

When two statements are joined by AND to produce a compound statement, we need a way of deciding whether the compound statement is true or false, based on the truth or falsity of its component statements. Let's build these with an example. Define statements p and q as follows:

$$p: \text{ Meghan is at least 25 years old.}$$
$$q: \text{ Meghan has a valid driver's license.}$$

Now let's create the statement we call "p and q," which we write as

$$p \wedge q: \text{ Meghan is at least 25 years old, and she has a valid driver's license.}$$

If you know the truth or falsity of p and q individually, how would you be inclined to categorize $p \wedge q$?[1] Naturally, the only way that we would consider $p \wedge q$ to be true is if both p and q are true. In any other instance, we would say $p \wedge q$ is false. So whether $p \wedge q$ is in $\mathcal{T}$ or $\mathcal{F}$ depends on whether p and q are in $\mathcal{T}$ or $\mathcal{F}$ individually. We illustrate the results of all different combinations in Table 1.2. Notice how the truth table is constructed with four rows, systematically displaying all possible combinations of T and F for p and q.

p	q	$p \wedge q$
T	T	T
T	F	F
F	T	F
F	F	F

(1.2)

Definition 1.1.3 Given two statements p and q, we define the statement $p \wedge q$ (p and q) to be true precisely when both p and q are true, as illustrated in Table 1.2.

1.1.4 Combining Statements with OR

Define statements p and q by

$$p: \text{ Meghan has insurance that covers her for any car she drives.}$$
$$q: \text{ Meghan bought the optional insurance provided by the car rental company.}$$

[1] Pretend Meghan is standing at a car rental counter and must answer yes or no to the question "Are you at least 25 years old and have a valid driver's license?"

The compound statement we call "p or q" is written

$p \vee q$: Meghan has insurance that covers her for any car she drives, or she bought the optional insurance provided by the car rental company.

How should we assign T or F to $p \vee q$ based on the truth or falsity of p and q individually?[2] We define $p \vee q$ to be true if *at least one* of the two statements is true. See Table 1.3.

p	q	$p \vee q$
T	T	T
T	F	T
F	T	T
F	F	F

(1.3)

Definition 1.1.4 Given two statements p and q, we define the statement $p \vee q$ (p or q) to be true precisely when at least one of p and q is true, as illustrated in Table 1.3.

Our conversational language is sometimes ambiguous when we use the word OR, and we often use OR in ways that are different from the mathematical use. For example, suppose a friend asks you "Did you have a cheeseburger or pizza for lunch?" In conversation, you would answer, "Oh, I had pizza," and your friend would likely conclude that you did not have the cheeseburger. In mathematics, however, you should answer your friend's question with a simple yes or no, depending on whether you had at least one of these two items.

To address a statement involving p and q that is true when precisely one of them is true, we use the *exclusive or*, which is written $p \dot{\vee} q$. See Table 1.4.

p	q	$p \dot{\vee} q$
T	T	F
T	F	T
F	T	T
F	F	F

(1.4)

Now we can build all kinds of compound statements.

[2] Pretend Meghan's friend is worried about being covered in case of an accident. He asks Meghan, "Do you have your own insurance, or did you buy the optional coverage provided by the rental company?" Under what circumstances should she say yes?

Example 1.1.5 Construct a truth table for the statement $(p \wedge q) \vee (\neg p \wedge \neg q)$. See Table 1.5.

Solution

p	q	$\neg p$	$\neg q$	$p \wedge q$	$\neg p \wedge \neg q$	$(p \wedge q) \vee (\neg p \wedge \neg q)$
T	T	F	F	T	F	T
T	F	F	T	F	F	F
F	T	T	F	F	F	F
F	F	T	T	F	T	T

(1.5)

■

Exercise 1.1.6 Use the layout of Table 1.6 to construct truth table columns for $p \wedge (q \vee r)$ and $(p \wedge q) \vee (p \wedge r)$.

p	q	r	$q \vee r$	$p \wedge (q \vee r)$	$p \wedge q$	$p \wedge r$	$(p \wedge q) \vee (p \wedge r)$
T	T	T					
T	T	F					
T	F	T					
T	F	F					
F	T	T					
F	T	F					
F	F	T					
F	F	F					

(1.6)

Exercise 1.1.7 Construct truth tables for the following statements.

(a) $p \vee (q \vee r)$

(b) $\neg(p \wedge q) \vee (p \vee q)$

(c) $(p \vee q) \vee r$

(d) $\neg p \vee q$

(e) $p \wedge (q \vee \neg p)$

1.1.5 Logical Equivalence

There are often several equivalent ways to say the same thing. We need to address the situation where two different constructs involving statements should be interpreted as having the same meaning, or as being *logically equivalent*. As a trivial example, consider that $p \wedge q$ should certainly be viewed as having the same meaning as $q \wedge p$. To build a truth table would produce identical columns for $p \wedge q$ and $q \wedge p$. This is the way we define *logical equivalence*.

Definition 1.1.8 Two statements are said to be *logically equivalent* if they have precisely the same truth table values. If U and V are logically equivalent, we write $U \Leftrightarrow V$.

Notice that the two statements in Exercise 1.1.6 are logically equivalent. To say "p AND either q or r" has the same meaning to us as "p and q, OR p and r." This is a sort of distributive property, where $\wedge$ distributes over $\vee$, just like multiplication distributes over addition in the real numbers.

Exercise 1.1.9 Show that $\vee$ distributes over $\wedge$ by showing that $p \vee (q \wedge r)$ is logically equivalent to $(p \vee q) \wedge (q \vee r)$.

Parts (a) and (b) of Exercise 1.1.7 show that $(p \vee q) \vee r$ is logically equivalent to $p \vee (q \vee r)$. We say that $\vee$ has the associative property, and we may allow ourselves the freedom to write $p \vee q \vee r$ to mean either $(p \vee q) \vee r$ or $p \vee (q \vee r)$.

Exercise 1.1.10 Does $\wedge$ have the associative property? Verify your answer with a truth table.

Exercise 1.1.11 Construct a statement using only p, q, $\wedge$, $\vee$, and $\neg$ that is logically equivalent to $p \dot{\vee} q$. Demonstrate logical equivalence with a truth table.

Exercise 1.1.12 Below are two logical equivalences called *DeMorgan's laws* (a name you will want to remember). Verify these forms of DeMorgan's laws with truth tables.

(a) $\neg(p \wedge q) \Leftrightarrow \neg p \vee \neg q$

(b) $\neg(p \vee q) \Leftrightarrow \neg p \wedge \neg q$

With Exercises 1.1.7(a) and (c), 1.1.10, and 1.1.12(a), we can extend one form of DeMorgan's law by using the following manipulation of logical symbols:

$$\begin{aligned}\neg(p \wedge q \wedge r) &\Leftrightarrow \neg[(p \wedge q) \wedge r] \Leftrightarrow \neg(p \wedge q) \vee \neg r \\ &\Leftrightarrow (\neg p \vee \neg q) \vee \neg r \Leftrightarrow \neg p \vee \neg q \vee \neg r\end{aligned} \tag{1.7}$$

Exercise 1.1.13 Mimic the manipulation in (1.7) to show that $\neg(p \vee q \vee r)$ is logically equivalent to $\neg p \wedge \neg q \wedge \neg r$.

Exercise 1.1.14 Use DeMorgan's laws and symbolic manipulation to show that $\neg[(p \vee q) \wedge r]$ is logically equivalent to $(\neg p \wedge \neg q) \vee \neg r$.

1.1.6 Tautologies and Contradictions

Sometimes a truth table column just happens to have TRUE values all the way down, as in Exercise 1.1.7(b). Another easy example is $p \vee \neg p$. A statement like "Either Meghan has a valid driver's license, or she does not" would make you think, "Of course!" or "Naturally this is a true statement regardless of the circumstances." A statement whose truth table values are all TRUE is called a *tautology*.

The negation of a tautology is called a *contradiction*, and the truth table values of a contradition are all FALSE. In the same way that a tautology is the kind of statement that makes you think "Of course!," a contradiction makes you think "No way can that ever be true!" A really easy example of a contradiction is $p \wedge \neg p$. Since p and $\neg p$ cannot ever both be true, $p \wedge \neg p$ is always false. Tautologies and contradictions are very useful, as we will begin to see in Chapter 2.

EXERCISE 1.1.15 Show that the statement $(p \wedge q) \vee (\neg p \vee \neg q)$ is a tautology.

1.2 If-Then Statements

In this section, we want to do two things: (1) Consider the logical structure of the statement that p *implies* or *necessitates* q and its variations; and (2) return to the idea of logical equivalence and its connection to tautologies.

1.2.1 If-Then Statements Defined

On the first day of class, your professor makes the following statement:

Either you turn in your homework on time or you get a zero on it.

Being the quick-witted logician, you immediately make mental definitions of the following statements:

p : Your homework is late.

q : You get a zero on your homework.

and you note that your professor's statement becomes $\neg p \vee q$. You also recall from Exercise 1.1.7(d) that the truth table for $\neg p \vee q$ has only one FALSE entry, and that is the entry where p is true but q is false. You think, "Late homework guarantees a zero, but even if I turn it in on time, I still might get a zero." Furthermore, you observe that it is natural to rephrase your professor's statement as an if-then statement in the following way:

If your homework is late, then you get a zero on it.

This statement is called "If p, then q" or simply "p implies q," and is written $p \to q$.

Definition 1.2.1 The statement $p \to q$ is defined to be logically equivalent to $\neg p \vee q$, as illustrated in Table 1.8. We call p the *hypothesis* and q the *conclusion*.

p	q	$p \to q$
T	T	T
T	F	F
F	T	T
F	F	T

(1.8)

Example 1.2.2 Construct truth tables for the following statements.

1. $p \to \neg q$
2. $(p \wedge q) \to r$

Solution Remember, the only FALSE entry in an if-then column occurs when the hypothesis is true and the conclusion is false. See Tables 1.9 and 1.10.

p	q	$\neg q$	$p \to \neg q$
T	T	F	F
T	F	T	T
F	T	F	T
F	F	T	T

(1.9)

p	q	r	$p \wedge q$	$(p \wedge q) \to r$
T	T	T	T	T
T	T	F	T	F
T	F	T	F	T
T	F	F	F	T
F	T	T	F	T
F	T	F	F	T
F	F	T	F	T
F	F	F	F	T

(1.10)

■

EXERCISE 1.2.3 Construct truth tables for the following statements.

(a) $\neg p \to \neg q$

(b) $q \to p$

(c) $\neg q \to \neg p$

(d) $(p \vee q) \to r$

(e) $(p \to q) \wedge (q \to p)$

(f) $p \to (q \vee r)$

EXERCISE 1.2.4 Verify that each of the following statements is a tautology.

(a) $[(p \to q) \wedge p] \to q$ (modus ponens)

(b) $[(p \to q) \wedge \neg q] \to \neg p$ (modus tollens)

(c) $[(p \vee q) \wedge \neg q] \to p$ (disjunctive syllogism)

(d) $[(p \to q) \wedge (q \to r)] \to (p \to r)$ (hypothetical syllogism)

(e) $\{[(p \wedge q) \to r] \wedge p\} \to (q \to r)$

Example 1.2.5 Rephrase the following statements in the form "If p, then q."

1. Every senior in the class made an A on the final exam.
2. An integer is prime only if it is greater than 1.

Solution

1. If you are a senior in the class, then you made an A on the final exam.
2. If an integer is prime, then it is greater than 1. ■

EXERCISE 1.2.6 Rephrase the following statements in the form "If p, then q."

(a) Either $a > 0$ or $a \leq 0$.

(b) Either I'm crazy, or there is a pink elephant floating overhead.

(c) Either it's hot in here, or I'm coming down with the flu.

(d) Every time a bell rings, an angel gets his wings.

(e) Only fools fall in love.

(f) Only koalas eat eucalyptus leaves.

(g) It only hurts when I laugh.[3]

[3] Is this saying that all laughter is painful?

(h) The only solutions to the equation $x^2 - x = 0$ are nonnegative.

(i) Drivers use the carpool lane only if there are at least two people in the car.

(j) A function is differentiable only if it is continuous.

(k) U only if V.

(l) Brett won't babysit for us unless we pay him minimum wage.

(m) Unless you follow the instructions, the cake won't turn out right.

(n) No shoes, no shirt, no service.

EXERCISE 1.2.7 The following statements are either in the form $p \to q$ or can be stated in that form. Rephrase them in the equivalent form $\neg p \vee q$.

(a) If you don't cross at the crosswalk, you'll get run over.

(b) If my boss is telling me the truth, then I will get a bonus this year.

(c) If you don't pay your bill, they will shut off your electricity.

(d) Unless you follow the instructions, the cake won't turn out right.

(e) Brett won't babysit for us unless we pay him minimum wage.

(f) A function is differentiable only if it is continuous.

1.2.2 Variations on $p \to q$

Given two statements p and q, we might want to analyze other possible if-then statements besides $p \to q$. Let's define the following statements:

p : Rudy has calculus homework.

q : Rudy goes to the library.

Letting $p \to q$ be our primary statement, we can put p and q together in other important ways.

$p \to q$	If Rudy has calculus homework, then he goes to the library.	(Primary statement)
$q \to p$	If Rudy goes to the library, then he has calculus homework.	(Converse)
$\neg p \to \neg q$	If Rudy does not have calculus homework, then he does not go to the library.	(Inverse)
$\neg q \to \neg p$	If Rudy does not to go the library, then he does not have calculus homework.	(Contrapositive)

Example 1.2.8 In Exercise 1.2.3, you constructed truth tables for $q \to p$, $\neg p \to \neg q$, and $\neg q \to \neg p$. The truth table values are displayed in Table 1.11. Notice that $p \to q$ is logically equivalent to its contrapositive, and its converse and inverse are logically equivalent.

p	q	$\neg p$	$\neg q$	$p \to q$	$q \to p$	$\neg p \to \neg q$	$\neg q \to \neg p$
T	T	F	F	T	T	T	T
T	F	F	T	F	T	T	F
F	T	T	F	T	F	F	T
F	F	T	T	T	T	T	T

(1.11)

■

Exercise 1.2.9 Suppose x represents some specific, but unknown real number. State the converse, inverse, and contrapositive of the following statement:

$$\text{If } x > 1, \text{then } x^3 - x > 0.$$

An important construct involving if-then is the statement that is true precisely when p and q are either both true or both false. For example, if an integer is divisible by 6, then it is divisible by both 2 and 3. Conversely, if an integer is divisible by both 2 and 3, then it is divisible by 6. We say that an integer is divisible by 6 *if and only if* it is divisible by both 2 and 3.

Definition 1.2.10 Given statements p and q, the statement $p \leftrightarrow q$ (read "p if and only if q," and often written "p iff q") is defined to be true precisely when p and q are either both true or both false.

Exercise 1.2.11 Construct a truth table for $(p \to q) \wedge (q \to p)$ to show that it is logically equivalent to $p \leftrightarrow q$.

Exercise 1.2.12 Let p, q, and r represent, respectively, the statements that Penelope, Quentin, and Rhonda studied for the math test. Using $p, q, r, \neg, \wedge, \vee, \dot{\vee}$, $\to$, and $\leftrightarrow$, translate the following sentences into a symbolic logical construction.

(a) Rhonda studied for the test, but Penelope did not.

(b) If Rhonda studied for the test, then so did Penelope.

(c) If Rhonda did not study for the test, then Penelope and Quentin didn't either.

(d) Quentin and Rhonda either both studied or both did not study.

(e) Neither Penelope nor Rhonda studied for the test.

(f) If Penelope did not study for the test, then neither did Rhonda.

(g) If either Penelope or Rhonda studied for the test, then so did Quentin.

(h) Precisely one of Penelope and Quentin studied for the test.

Exercise 1.2.13 Which two of the statements in Exercise 1.2.12 are logically equivalent to each other?

1.2.3 Logical Equivalence and Tautologies

In Section 1.1, we defined two statements to be logically equivalent if they have exactly the same truth table values. With the definition of $p \leftrightarrow q$, we can now offer an alternate definition of logical equivalence.

Definition 1.2.14 Two statements U and V are logically equivalent if the statement $U \leftrightarrow V$ is a tautology.

As we noted in Section 1.1.5, the significance of statements being logically equivalent is that they are different ways of saying precisely the same thing. If U is logically equivalent to V, then knowing U is true guarantees that V is true, and vice versa.

If $U \leftrightarrow V$ is a tautology, then $U \rightarrow V$ and $V \rightarrow U$ must both be tautologies, too. Loosely speaking, the truth of U is sufficiently strong to imply the truth of V and vice versa.

Example 1.2.15 Let U be the statement $p \rightarrow q$ and let V be the statement $\neg q \rightarrow \neg p$. Show that U and V are logically equivalent using Definition 1.2.14.

Solution We already know that the truth table columns for $p \rightarrow q$ and $\neg q \rightarrow \neg p$ are identical, but we are asked to use Definition 1.2.14. So we construct $(p \rightarrow q) \leftrightarrow (\neg q \rightarrow \neg p)$ and show that it is a tautology. The details are displayed in Table 1.12.

p	q	$\neg p$	$\neg q$	U	V	$U \rightarrow V$	$V \rightarrow U$	$(U \rightarrow V) \wedge (V \rightarrow U)$
T	T	F	F	T	T	T	T	T
T	F	F	T	F	F	T	T	T
F	T	T	F	T	T	T	T	T
F	F	T	T	T	T	T	T	T

(1.12)

Since the last column of Table 1.12 is a tautology, U and V are logically equivalent. Notice that the columns $U \rightarrow V$ and $V \rightarrow U$ are tautologies to make the last column a tautology. ■

Now let's consider the situation where $U \leftrightarrow V$ is not a tautology, but one of $U \rightarrow V$ or $V \rightarrow U$ is. If $U \rightarrow V$ is a tautology while $V \rightarrow U$ is not, then we say that

U is a *stronger statement* than V. This means that the truth of U necessitates the truth of V, but the truth of V is not necessarily accompanied by the truth of U.

Example 1.2.16 Which statement is stronger, p or $p \wedge q$? Verify with a truth table.

Solution See Table 1.13. Since $(p \wedge q) \rightarrow p$ is a tautology while $p \rightarrow (p \wedge q)$ is not, $p \wedge q$ is stronger than p. Knowing $p \wedge q$ is true guarantees that p is true, but knowing p is true does not guarantee that $p \wedge q$ is true.

p	q	$p \wedge q$	$(p \wedge q) \rightarrow p$	$p \rightarrow (p \wedge q)$
T	T	T	T	T
T	F	F	T	F
F	T	F	T	T
F	F	F	T	T

(1.13)

■

Example 1.2.17 Which statement do you think is stronger, $(p \wedge q) \rightarrow r$ or $p \rightarrow r$? Determine for sure with a truth table.

Solution Writing $(p \wedge q) \rightarrow r$ as U and $p \rightarrow r$ as V, we need truth table values for $U \rightarrow V$ and $V \rightarrow U$. (See Table 1.14.) Since $V \rightarrow U$ is a tautology and $U \rightarrow V$ is not, V is stronger than U.

p	q	r	$p \wedge q$	$U : (p \wedge q) \rightarrow r$	$V : p \rightarrow r$	$U \rightarrow V$	$V \rightarrow U$
T	T	T	T	T	T	T	T
T	T	F	T	F	F	T	T
T	F	T	F	T	T	T	T
T	F	F	F	T	F	F	T
F	T	T	F	T	T	T	T
F	T	F	F	T	T	T	T
F	F	T	F	T	T	T	T
F	F	F	F	T	T	T	T

(1.14)

■

Notice this important fact. Statements U and V from Example 1.2.17 have the same conclusion, but the hypothesis condition for U is *stronger* than the hypothesis condition for V. This makes U a *weaker* statement than V.

What is the significance of one statement being stronger than another? Here is an example to illustrate Example 1.2.17. Define the following statements.

p : Meghan is at least 25 years old.

q : Meghan has a valid driver's license.

r : Meghan is allowed to rent a car.

What does it mean to say that $p \to r$ is stronger than $(p \wedge q) \to r$? The statement $(p \wedge q) \to r$ says that age and a license will guarantee your eligibility to rent a car, while $p \to r$ says that age alone is sufficient to be eligible. So if everyone at least 25 years old can rent a car, then certainly everyone at least 25 with a license can, too. Thus if $p \to r$ is true, so is $(p \wedge q) \to r$. On the other hand, just because licensed people at least 25 years old can rent a car, it does not follow that all people over 25 can do the same. That is, the truth of $(p \wedge q) \to r$ does not imply that $p \to r$ is true.

EXERCISE 1.2.18 For each of the following pairs of statements, use a truth table to determine whether they are logically equivalent, or that one is stronger than the other.

(a)	p	$p \vee q$
(b)	$p \to r$	$(p \vee q) \to r$
(c)	$p \to r$	$(p \to q) \wedge (q \to r)$
(d)	$(p \vee q) \to r$	$(p \to r) \vee (q \to r)$
(e)	$(p \vee q) \to r$	$(p \to r) \wedge (q \to r)$
(f)	$p \to (q \wedge r)$	$(p \to q) \wedge (p \to r)$
(g)	$p \dot{\vee} q$	$(p \vee q) \wedge \neg q$
(h)	$p \to (q \vee r)$	$(p \wedge \neg q) \to r$
(i)	$(p \wedge q) \to r$	$(p \wedge \neg r) \to \neg q$
(j)	$(p \leftrightarrow q) \wedge (r \leftrightarrow s)$	$(p \vee r) \leftrightarrow (q \vee s)$
(k)	$p \leftrightarrow q$	$\neg p \leftrightarrow \neg q$

EXERCISE 1.2.19 This exercise investigates the general principle that an implication statement is strengthened when its hypothesis condition is weakened.

(a) Suppose U and V are two statements, and U is stronger than V. What must be true about the truth table entries for U as compared to those for V?[4]

(b) If U is stronger than V, why does that make $U \to r$ weaker than $U \to r$?[5]

[4] The set of F entries for one must be a proper subset of the F entries for the other.

[5] Use your answer to part (a), and the definition of $p \to q$.

(c) Which of the following pairs of statements is stronger? Explain without using a truth table.

(i) $p \to r$ and $(p \vee q) \to r$

(ii) $[p \wedge (q \to s)] \to t$ and $\{p \wedge [(q \wedge r) \to s]\} \to t$

Example 1.2.20 Without justifying by proof, state whether statements in each of the following pairs are logically equivalent, whether one is stronger than the other, or neither.

1.

$p:$ x is an integer that is divisible by 6.

$q:$ x is an integer that is divisible by 3.

2.

$$p: \; x^3 - 4x^2 + 4x = 0$$

$$q: \; x \in \{0, 2, 4\}$$

3.

$$p: \; -2 < x < 2$$

$$q: \; x^2 < 4$$

4.

$p:$ x is an integer multiple of 10.

$q:$ x is an integer multiple of 15.

Solution

1. If an integer is divisible by 6, then it is divisible by both 2 and 3. However, divisibility by 3 does not necessitate divisibility by 6. For example, 9 is divisible by 3 but not divisible by 6. So p is stronger than q.
2. Factoring and solving the equation in p yield $x = 0$ or $x = 2$. So if $x = 0$ or $x = 2$, then $x \in \{0, 2, 4\}$. But $x = 4$ is not a solution to the equation $x^3 - 4x^2 + 4x = 0$. Thus if $x \in \{0, 2, 4\}$, it does not follow that $x^3 - 4x^2 + 4x = 0$. Therefore p is stronger than q.
3. A particular value of x will either make p and q both true or both false. So they are logically equivalent.
4. If x is an integer multiple of 10, it might or might not be an integer multiple of 15, as is illustrated by $x = 20$ and $x = 30$. Furthermore, just because x is an

integer multiple of 15, it might or might not be an integer multiple of 10, as is illustrated by $x = 30$ and $x = 45$. Thus neither of these statements is stronger than the other. ■

EXERCISE 1.2.21 Suppose N is some unknown real number and f is a given function. Without justifying by proof, state whether statements in each of the following pairs are logically equivalent, whether one is stronger than the other, or neither.

(a)

N is a natural number.

N is an integer.

(b)

$$N > 5$$

$$N > 10$$

(c)

If $N > 5$, then $f(N) > 0$.

If $N > 10$, then $f(N) > 0$.

(d)

If $N > 5$ and N is an integer, then $f(N) > 0$.

If $N > 10$, then $f(N) > 0$.

1.3 Universal and Existential Quantifiers

Most of the examples of if-then statements in Section 1.2 were not about mathematical objects, but about people, situations, and behaviors. The reason is that mathematical if-then statements generally involve indeterminates (as the statements defined on p. 21), and these are inextricably tied to the presence of what we call universal and existential quantifiers.

We might make a statement like:

All squares are rectangles,

which we might phrase as

If x is a square, then x is a rectangle,

but which a more formal logician would express as

For all x, if x is a square, then x is a rectangle.

This last statement might sound odd, but you will see in Section 1.4 why the presence of "for all x" is important. In this section, we will express statements in a variety of ways, each of which is natural and appropriate in a certain context.

1.3.1 The Universal Quantifier

The inequality $n^2 < 2^n$ is true for all natural numbers $n \geq 5$. (You will prove this in Section 3.4.) We could express this in any of the following ways.

For all natural numbers n, if $n \geq 5$, then $n^2 < 2^n$.

For all n, if $n \in \mathbb{N}$ and $n \geq 5$, then $n^2 < 2^n$.

$(\forall n)[(n \in \mathbb{N} \wedge n \geq 5) \rightarrow (n^2 < 2^n)]$.

These statements are arranged in increasing order of formality. The last one uses the symbol $\forall$ (for all), which is called the *universal quantifier*.

Example 1.3.1 The following are all acceptable ways of making the same universal statement.

Every element of the set B is negative.

For all $x \in B$, $x < 0$.

$(\forall x \in B)(x < 0)$.

For all x, if $x \in B$, then $x < 0$.

$(\forall x)(x \in B \rightarrow x < 0)$. ■

All five of the statements in Example 1.3.1 have their place in mathematical discourse, from the very informal first one to the very formal last one. Rather than think of the last statement as a formalization of the first, however, it is most helpful to think of the first as a more conversational equivalent of the last.

Example 1.3.2 Suppose f is a given function. Then the following statements are equivalent.

The graph of f does not intersect the x-axis.

For all real numbers x, $f(x) \neq 0$.

$(\forall x \in \mathbb{R})(f(x) \neq 0)$.

For all x, if $x \in \mathbb{R}$, then $f(x) \neq 0$.

$(\forall x)(x \in \mathbb{R} \rightarrow f(x) \neq 0)$.

The most general form for a statement involving the universal quantifer would look something like this. If $P(x)$ is some property stated in terms of x, such as "$x \in \mathbb{R} \rightarrow f(x) \neq 0$," then a general statement involving the universal quantifier would be written

$$(\forall x)(P(x)) \tag{1.15}$$

and would be read, "For all x, $P(x)$." ■

Exercise 1.3.3 Let D be the set of all differentiable functions and C the set of all continuous functions. Construct several statements (as in Examples 1.3.1 and 1.3.2) that use the universal quantifier and are equivalent to the statement

If a function is differentiable, then it is continuous.

Exercise 1.3.4 Construct several statements that use the universal quantifier and are equivalent to the statement

The only real number solutions to the equation $x^2 - x = 0$ are nonnegative.

1.3.2 The Existential Quantifier

The inequality $n^2 < 2^n$ is true for $n \geq 5$, but not if $1 \leq n \leq 4$. In other words, there exist natural numbers n (at least one such number) such that $n^2 \geq 2^n$.

There exists a natural number n such that $n^2 \geq 2^n$.

There exists n such that $n \in \mathbb{N}$ and $n^2 \geq 2^n$.

$(\exists n)(n \in \mathbb{N} \wedge n^2 \geq 2^n)$.

The expression "there exists," written mathematically as $\exists$, is called the existential quantifier.

Example 1.3.5 The following are all acceptable ways of making the same existential statement.

Some element of the set B is positive.

There exists $x \in B$ such that $x > 0$.

$(\exists x \in B)(x > 0)$.

There exists x such that $x \in B$ and $x > 0$.

$(\exists x)(x \in B \wedge x > 0)$.

The most general form of an existential statement involving a property $P(x)$ would be written

$$(\exists x)(P(x)) \tag{1.16}$$

and would be read, "There exists x such that $P(x)$." ■

Exercise 1.3.6 Let D be the set of all differentiable functions and C the set of all continuous functions. Construct several statements (as in Example 1.3.5) that use the existential quantifier and are equivalent to the statement

Some continuous functions are not differentiable.

Exercise 1.3.7 Construct several statements that use the existential quantifier and are equivalent to the statement

The equation $x^2 + x = 1$ has a negative real number solution.

Now the plot thickens.

Example 1.3.8 Suppose F is a set of functions. Here are three equivalent statements.

The graph of every function in F intersects the x-axis at least once.

For all $f \in F$, there exists $x \in \mathbb{R}$ such that $f(x) = 0$.

$(\forall f \in F)(\exists x \in \mathbb{R})(f(x) = 0)$. ■

Exercise 1.3.9 Let F be a set of functions. Restate the following using universal and existential quantifiers.

(a) There is a function in F whose graph intersects the x-axis.

(b) No function in F has a graph that intersects the x-axis.

(c) There is a function in F whose graph does not intersect the x-axis.

Sometimes a universal statement might be expressed informally by placing the $\forall$ phrase *after* the property that is universally true. This is merely a way to make a sentence sound more natural. We would not write the statement that way symbolically. The next example illustrates this with a range of expressions from informal to formal.

Example 1.3.10 The following statements are equivalent.

> There is an element of set A that is less than every element of set B.
>
> There exists $x \in A$ such that $x < y$ for all $y \in B$.
>
> There exists $x \in A$ such that, for all $y \in B$, $x < y$.
>
> $(\exists x \in A)(\forall y \in B)(x < y)$. ■

EXERCISE 1.3.11 For each of the following statements, use your creativity and construct several statements that are logically equivalent to the given statement.

(a) No element of the set A exceeds m.

(b) Some element of the set A exceeds m.

(c) A contains an element that is greater than every element of B.

(d) Every element of A is greater than every element of B.

(e) There is some element of A that is greater than some element of B.

(f) Every element of A is greater than some element of B.

(g) There is no real number that is in both sets A and B.

(h) Every element of the set A can be written as the sum of the squares of two integers.

(i) The equation $x^2 + 1 = 0$ has no real number solution.

(j) Every real number has an additive inverse.

(k) If n is an odd integer, then n^2 is an odd integer.

(l) There is a real number that does not have a multiplicative inverse.

1.3.3 Unique Existence

Sometimes it is important to know not only that something exists, but also that exactly one such thing exists. If there exists exactly one thing with a certain property, we say that it exists *uniquely*. The mathematical statement

$$(\exists! x)(P(x)) \tag{1.17}$$

is read, "There exists a unique x such that $P(x)$."

To define unique existence, we must amend the statement $(\exists x)(P(x))$ to include the additional stipulation that there is no more than one such x. The standard way of defining unique existence is the following.

Definition 1.3.12 We say that there exists a unique x with property P provided two things are true. First, there exists x with property P, and second, for all x_1 and x_2, if x_1 and x_2 both have property P, then x_1 and x_2 are the same. In other words, the logical statement $(\exists! x)(P(x))$ is defined by

$$[(\exists x)(P(x))] \wedge [(\forall x_1)(\forall x_2)(P(x_1) \wedge P(x_2)) \rightarrow (x_1 = x_2)] \tag{1.18}$$

Unique existence is really a compound AND statement. The first component says that at least one thing with the property exists, and the second component says that at most one thing with the property exists.

Example 1.3.13 Reword the following statements of unique existence in a form like that in Definition 1.3.12.

1. There exists a unique real number x such that $x^3 = 8$.
2. For all $a \in \mathbb{R}$, there exists a unique $b \in \mathbb{R}$ such that $a + b = 0$.

Solution

1. There exists a real number x such that $x^3 = 8$, and for all $x_1, x_2 \in \mathbb{R}$, if $x_1^3 = 8$ and $x_2^3 = 8$, then $x_1 = x_2$.
2. For all $a \in \mathbb{R}$, there exists $b \in \mathbb{R}$ such that $a + b = 0$, and for all $b_1, b_2 \in \mathbb{R}$, if $a + b_1 = 0$ and $a + b_2 = 0$, then $b_1 = b_2$.

When we write these unique existence statements in prose, we can give ourselves the freedom to omit the universal quantifier phrase, knowing it is understood. So we could reword these slightly as

1. There exists a real number x such that $x^3 = 8$, and if $x_1, x_2 \in \mathbb{R}$ satisfy $x_1^3 = 8$ and $x_2^3 = 8$, then $x_1 = x_2$.

2. For all $a \in \mathbb{R}$, there exists $b \in \mathbb{R}$ such that $a + b = 0$, and if $b_1, b_2 \in \mathbb{R}$ satisfy $a + b_1 = 0$ and $a + b_2 = 0$, then $b_1 = b_2$. ■

EXERCISE 1.3.14 Reword each of the following unique existence statements in a form like that in Definition 1.3.12.

(a) The set A contains precisely one element.

(b) The function f crosses the x-axis exactly once.

(c) There is a unique function f such that $f(0) = 1$ and $f'(x) = f(x)$ for all $x \in \mathbb{R}$.

(d) There exists a unique natural number n such that $n > 1$ and n is a factor of p.

(e) The function $f(x) = x^3 - x - 1$ crosses the x-axis precisely once in the interval $0 \le x \le 2$.

(f) The curves $x^2 + y^2 = 1$ and $15y = (x - 11/5)^2$ intersect at exactly one point in the xy-plane.

1.4 Negations of Statements

The defining characteristic of the negation of a statement is that the truth table values are exactly the opposite. In this section, we see how to construct negations of compound statements. Then, if someone makes an ugly statement like

$$(p \wedge q) \to (r \vee s)$$

and we want to say to that person, "Nope, you're wrong," then we'll have a way of writing

$$\neg[(p \wedge q) \to (r \vee s)]$$

in a more useful and understandable form.

1.4.1 Negations of AND and OR Statements

In Exercise 1.1.12, you showed that $\neg(p \wedge q)$ is logically equivalent to $\neg p \vee \neg q$, and that $\neg(p \vee q)$ is logically equivalent to $\neg p \wedge \neg q$. These facts are called *DeMorgan's laws*, and they provide us a way of expressing the negations of $\wedge$ and $\vee$ compound statements.

Example 1.4.1 Construct a negation of the following statements.

1. $p \wedge q \wedge r$
2. $(p \vee q) \wedge (r \vee s)$

Solution Applying DeMorgan's laws, we have

1. $\neg(p \wedge q \wedge r) \Leftrightarrow \neg p \vee \neg q \vee \neg r$
2. $\neg[(p \vee q) \wedge (r \vee s)] \Leftrightarrow \neg(p \vee q) \vee \neg(r \vee s) \Leftrightarrow (\neg p \wedge \neg q) \vee (\neg r \wedge \neg s)$ ■

Example 1.4.2 Use DeMorgan's laws to express in words a negation of the following statements.

1. Jacob has brown hair and blue eyes.
2. Either I'm crazy or there is a pink elephant floating overhead.
3. Meghan is at least 25 years old, she has a valid driver's license, and she either has her own insurance or has purchased coverage from the car rental company.

Solution Don't hesitate to word the statements in a way that makes them easy to understand.

1. Either Jacob does not have brown hair, or he does not have blue eyes.
2. I am not crazy, and there is no pink elephant floating overhead.
3. Either Meghan is under 25, or she doesn't have a valid driver's license, or she has no insurance of her own and has not purchased coverage from the car rental company. ■

EXERCISE 1.4.3 State in words a negation for each of the following statements.

(a) Either Melanie is naturally blond, or she bleaches her hair.

(b) Eric has a boarding pass and either a driver's license or passport.

(c) Either f is not continuous or it crosses the x-axis at some point. (f is a given function.)

(d) Been there; done that.

1.4.2 Negations of If-Then Statements

In Section 1.2, we defined $p \to q$ to be logically equivalent to $\neg p \vee q$. To construct a negation of $p \to q$, we can use this fact with DeMorgan's law.

$$\neg(p \to q) \Leftrightarrow \neg(\neg p \vee q) \Leftrightarrow (\neg\neg p) \wedge \neg q \Leftrightarrow p \wedge \neg q$$

This might be a little confusing at first, but later you will want to think of it in the following way. If someone makes a claim that p implies q, then he or she is claiming that the truth of p is by necessity accompanied by the truth of q. If you want to negate that claim, then you are suggesting that p can be true, while q is false. The techniques that follow for negating universal and existential statements will help clear that up.

Example 1.4.4 Construct a negation of the following statements.

1. $p \to (q \wedge r)$
2. $(p \to q) \vee (r \to s)$

Solution

1. $\neg[p \to (q \wedge r)] \Leftrightarrow p \wedge \neg(q \wedge r) \Leftrightarrow p \wedge (\neg q \vee \neg r)$
2. $\neg[(p \to q) \vee (r \to s)] \Leftrightarrow \neg(p \to q) \wedge \neg(r \to s) \Leftrightarrow (p \wedge \neg q) \wedge (r \wedge \neg s)$ ■

Example 1.4.5 State in words a negation of the following statements.

1. If Trent is a senior in the class, then he is exempt from the final exam.
2. If Meghan has rented a car, then she either has her own insurance or has bought coverage from the car rental company.

Solution

1. Trent is a senior in the class, and he is not exempt from the final exam.
2. Meghan rented a car, but she neither has her own insurance nor did she purchase coverage from the car rental company. ■

Exercise 1.4.6 Use the results of this section to construct the symbolic negation of the following statements.

(a) $p \wedge (q \vee r)$

(b) $p \to (q \vee r)$

(c) $(p \vee q) \to r$

(d) $(p \to q) \wedge (p \to r)$

(e) $[(p \wedge q) \to r] \to s$

(f) $p \leftrightarrow q$

Exercise 1.4.7 State in words a negation for each of the following statements.

(a) If Brett is on the guest list, then he will come to the party.

(b) If f is continuous, then it crosses the x-axis at some point. (f is a given function.)

(c) No shoes, no shirt, no service.

1.4.3 Negations of Statements with the Universal Quantifier

Suppose someone makes the following claim.

Every person in the class passed the first exam.

For you to deny this claim, what sort of statement would you make?[6] You would probably say something like

At least one person in the class failed the first exam.

This is where the most formal version of the universal statement comes in handy. Let C be the set of students in the class and let P be the set of students in the class who passed the first exam. We can write the original statement as

$$(\forall x)(x \in C \to x \in P) \tag{1.19}$$

and its negation would be

$$(\exists x)(x \in C \wedge x \notin P)$$

Notice in the negation that $(x \in C \wedge x \notin P)$ is a negation of the statement $(x \in C \to x \in P)$. Using the property notation $P(x)$, we can therefore say the following.

$$\neg[(\forall x)(P(x))] \Leftrightarrow (\exists x)(\neg P(x)) \tag{1.20}$$

So the trick to negating a universal statement is that the $\neg$ symbol crawls over the $(\forall x)$ and converts it to $(\exists x)$ as it goes.

Another symbolic rendering of statement (1.19) would be

$$(\forall x \in C)(x \in P)$$

whose negation would be

$$(\exists x \in C)(x \notin P)$$

Notice that the $(\forall x \in C)$ becomes $(\exists x \in C)$ as the $\neg$ symbol crawls over it, then the statement $(x \in P)$ is negated.

Example 1.4.8 State in words a negation of the following statements.

1. For all $x \in \mathbb{Z}$, $x^2 \geq 0$.
2. For all $x \in \mathbb{Z}$, if x is divisible by 6, then x is divisible by 3.
3. Every real number is either rational or irrational.

[6] One slacker is all it takes.

Solution

1. There exists $x \in \mathbb{Z}$ such that $x^2 < 0$.
2. There exists $x \in \mathbb{Z}$ such that x is divisible by 6 but not divisible by 3.
3. There exists a real number that is neither rational nor irrational. ■

Warning: Universal statements are often written in if-then form with the universal quantifier omitted. This is perfectly acceptable in mathematical discourse, but the presence of the universal quantifier must be understood if you are going to negate the statement correctly.

Example 1.4.9 A mathematical statement such as

$$\text{If } x > 1, \text{then } x^3 - x > 0.$$

is an acceptable simplification of the more formal

$$\text{For all } x, \text{if } x \in \mathbb{R} \text{ and } x > 1, \text{then } x^3 - x > 0.$$

even though it contains the indeterminate x and does not specify that the context is the real numbers. However, the formalized version must be understood in order to construct the correct negation:

$$\text{There exists a real number } x \text{ such that } x > 1 \text{ and } x^3 - x \leq 0. \quad ■$$

Exercise 1.4.10 State in words a negation for each of the following statements.

(a) If n is a natural number and $n \geq 5$, then $n^2 < 2^n$.

(b) The only solutions to the equation $x^2 - x = 0$ are nonnegative.

(c) On a clear day, you can see forever.

(d) If $-1 \leq x \leq 10$, then $x^2 \leq 100$.

(e) Every time a bell rings, an angel gets his wings.

(f) No one in the class made an A on the final.

1.4.4 Negations of Statements with the Existential Quantifier

Now we negate statements involving the existential quantifier. Consider the following.

Someone traveling in my car didn't chip in enough money for gas.

If we let C be the set of people riding in my car and T be the set of tightwads, the statement becomes

$$(\exists x)(x \in C \wedge x \in T)$$

To negate this, we avoid saying "no one" if possible. It is better to say what everyone did instead of what no one did. We could say

Everyone traveling in my car chipped in enough money for gas,

or

$$(\forall x \in C)(x \notin T)$$

which is equivalent to

$$(\forall x)(x \in C \to x \notin T)$$

Notice in the negation that $(x \in C \to x \notin T)$ is the negation of $(x \in C \wedge x \in T)$. Thus we arrive at the following.

$$\neg[(\exists x)(P(x))] \Leftrightarrow (\forall x)(\neg P(x)) \tag{1.21}$$

This looks similar to statement (1.20) in that negating an existential statement can be done by letting the $\neg$ symbol crawl over the $\exists$, changing it to $\forall$ as it goes.

If someone says to you that a certain something exists, and it is a claim you want to deny, then avoid using the expression "There does not exist." Instead, express your negation in positive language by saying that all things fail to have the desired property.

Example 1.4.11 Construct a negation of the following statements.

1. $(\exists x \in \mathbb{N})(x \leq 0)$
2. $(\forall \epsilon > 0)(\exists n \in \mathbb{N})(1/n < \epsilon)$

Solution

1. $(\forall x \in \mathbb{N})(x > 0)$
2. $(\exists \epsilon > 0)(\forall n \in \mathbb{N})(1/n \geq \epsilon)$ ■

Example 1.4.12 State in words a negation of the following statements.

1. Someone in this class cheated on the final exam.
2. There exists a natural number x such that $x \leq y$ for all $y \in \mathbb{N}$.
3. There exists $M \in \mathbb{R}$ such that $x \leq M$ for all $x \in S$.
4. For all $a \in A$ and $\epsilon > 0$, there exists $\delta > 0$ such that, if $0 < |x - a| < \delta$, then $|f(x) - L| < \epsilon$.

Solution

1. Everyone in this class did not cheat on the final exam.
2. For every natural number x, there exists a natural number y such that $x > y$.
3. For all real numbers M, there exists $x \in S$ such that $x > M$.
4. There exists $a \in A$ and $\epsilon > 0$ such that, for all $\delta > 0$, there exists an x satisfying $0 < |x - a| < \delta$ and $|f(x) - L| \geq \epsilon$. ■

EXERCISE 1.4.13 Construct a negation of the following statements.

(a) $a > 1 \wedge a \leq 10$

(b) $a \geq 120 \vee a \leq 80$

(c) $(\forall x)(x^2 - 4x + 3 = 0 \rightarrow x \geq 1)$

(d) $(\exists x)(x \geq 2 \wedge x^2 - 3x = 2)$

(e) $(\forall a \in \mathbb{R})(\exists b \in \mathbb{R})(a + b = 0)$

(f) $(\exists! x)(P(x))$

(g) $(\forall a \in \mathbb{R})(\exists! b \in \mathbb{R})(a + b = 0)$

(h) $(\exists m \in A)(\forall x \in A)(x \leq m)$

(i) $(\forall x)(x \in A \rightarrow x \in B)$

(j) $(\exists x)(x \in A \wedge x \in B)$

(k) There is an element of set A that is less than every element of set B.

(l) Every element of A is greater than every element of B.

(m) There exists an element of the set A that exceeds m.

(n) The graph of every function in F intersects the x-axis at least once.

(o) There is a function in F whose graph intersects the x-axis.

(p) No function in F has a graph that intersects the x-axis.

(q) There is a function in F whose graph does not intersect the x-axis.

(r) The equation $x^3 = 10$ has a real number solution.

(s) There exist $M_1, M_2 \in \mathbb{R}$ such that $M_1 \leq x \leq M_2$ for all $x \in S$.

(t) There exists a unique $x \in A$.

(u) Every real number has a multiplicative inverse.

(v) Every integer is either even or odd.

1.5 How We Write Proofs

A theorem is a statement of the form $U \to V$. To "prove" a theorem in the most abstract sense is to show that $U \to V$ is a tautology, for then we know that U is at least as strong as V. Thus, as we say, V follows from U, meaning that any context in which U is true will guarantee that V is also true.

Only the most theoretical mathematicians approach theorems as tautologies, and even then such mathematicians are probably more interested either in the philosophy of mathematical reasoning or in formal methods as it relates to fields such as computer science. In this text, we do hope that some mathematical philosophy rubs off on you. But we are primarily interested in proofs as almost all mathematicians write them. Thus we focus on what it means to write a proof in mathematical prose, so that it is accurate, clear, and readable. However, just so you will not feel cheated, we will provide one example (in Section 3.2) of a formalized theorem proved by demonstrating it is a tautology.

If we are not going to use tautologies in our proof writing, then why have we spent time studying logic and truth tables in this chapter? There are two primary reasons. One is that tautologies provide us with rules for valid reasoning that will undergird our mathematical writing. For example, in Exercise 1.2.4 you showed that $[(p \to q) \wedge p] \to q$ is a tautology. This tautology is called the *modus ponens,* and it is one of several formalized *rules of inference* that comprise mathematical reasoning. If we assume p and that $p \to q$, then the modus ponens rule of inference says we may therefore conclude q. It is not our purpose here to define all the rules of inference, though you have seen most of them at some point in this chapter. Exercise 1.2.4 contains the names of several rules of inference.

The second reason we have studied logic and truth tables is that it will provide us with approaches to the writing of proofs. In particular, we need to understand how to dissect mathematical statements (the term is to *parse* sentences) and see how they give order and structure to our proofs. We also need to know that we can sometimes write proofs by exploiting equivalent statements. Furthermore, a basic understanding of logic can immunize us from making certain *fallacies*—statements that are not tautologies and therefore not valid theorems.

In this section, we want to outline the four types of demonstrations you will employ throughout this text when you are asked to compose a proof of a mathematical proposition. Know that you will always use one of these methods of attack in any demonstration. Knowing which to choose depends on what your task is, and is illumined by experience.

1.5.1 Direct Proof

Most theorems are proved *directly*, which means that the statement of the theorem and its proof fit into the following template.

Theorem 1.5.1 (Sample). If p, then q.

Proof. Suppose p. Then .... Thus q. □

Our purpose here is only to see where a direct proof begins and ends. The "dot dot dot" is what varies from situation to situation. Here is an example theorem.

Theorem 1.5.2 (Sample). For all integers x and y, if x and y are both odd, then xy is odd.

Proof. Suppose x and y are both odd integers. Then .... Therefore xy is odd. □

Theorem 1.5.3 (Sample). If $a_n \to 0$ as $n \to \infty$ and $\langle b_n \rangle_{n=1}^{\infty}$ is bounded, then $a_n b_n \to 0$ as $n \to \infty$.

Proof. Suppose $a_n \to 0$ as $n \to \infty$ and that $\langle b_n \rangle_{n=1}^{\infty}$ is bounded. Then .... Thus $a_n b_n \to 0$ as $n \to \infty$. □

EXERCISE 1.5.4 Each of the following theorems is to be proved directly. State the first and last sentences of the proof of the theorem.

(a) If $f : A \to B$ and $g : B \to C$ are one-to-one, then $(g \circ f) : A \to C$ is one-to-one.

(b) If a set A contains all its cluster points, then A is closed.

1.5.2 Proof by Contrapositive

If your task is to prove a theorem of the form $p \to q$, then it suffices to prove any statement that is equivalent to $p \to q$. The contrapositive is the most common such equivalent, and a *proof by contrapositve* always goes something like this.

Theorem 1.5.5 (Sample). If p, then q.

Proof. Suppose $\neg q$. Then .... Thus $\neg p$. □

Theorem 1.5.6 (Sample). If $x \in \overline{A} \cup \overline{B}$, then $x \in \overline{A \cup B}$.

Proof. Suppose $x \notin \overline{A \cup B}$. Then .... Thus $x \notin \overline{A} \cup \overline{B}$. □

EXERCISE 1.5.7 Each of the following theorems is to be proved by contrapositive. State the first and last sentences of the proof of the theorem.

(a) If $A \subseteq B$, then $A - B$ is empty.

(b) If ab is irrational, then either a is irrational or b is irrational.

1.5.3 Proving a Logically Equivalent Statement

If a theorem says $p \to q$, then you may prove it by showing $\neg q \to \neg p$ because it is logically equivalent to the given theorem. There are many other mathematical propositions that are best proved by exploiting a logically equivalent form. Here is a common one.

Theorem 1.5.8 (Sample). $p \to (q \vee r)$.

Proof. Suppose $p \wedge \neg q$. Then .... Thus r. □

EXERCISE 1.5.9 Find the appropriate exercise in Section 1.2 that validates the proof of Theorem 1.5.8.

EXERCISE 1.5.10 Each of the following theorems might be proved naturally by exploiting a logical equivalence from Section 1.2. State the first and last sentences of the proofs of the theorem.

(a) $(p \wedge q) \to r$

(b) If p divides ab, then p divides a or p divides b.

(c) If a is rational and ab is irrational, then b is irrational.

1.5.4 Proof by Contradiction

If U represents the statement of a theorem, then U is a tautology. Thus $\neg U$ is a contradiction. Sometimes the most natural way to prove U is a tautology is to show that $\neg U$ is a contradiction. You're gonna love this.

Theorem 1.5.11 (Sample). U.

Proof. Suppose $\neg U$. Then ... KABOOM![7] This is a contradiction. Thus U. □

EXERCISE 1.5.12 Each of the following theorems is to be proved by contradiction. State the first and last sentences of the proofs of the theorem.

(a) $p \to q$

(b) The natural numbers are unbounded in the reals.

(c) There are infinitely many prime numbers.

(d) $1 > 0$

1.5.5 Disproving a Statement

Sometimes we are presented with a mathematical statement that is not true, and we are asked to verify this by *disproving* the given statement. To do this, we prove

[7] Mathematicians have been known to indulge in the occasional exclamation.

the negation of the statement. This is where our techniques of negation from Section 1.4 come in handy.

Example 1.5.13 Each of the following statements can be disproved. By negating the statement, provide the statement that could be proved true.

1. For all sets A, B, and C, if $A \cup C = B \cup C$, then $A = B$.
2. For all sets A, B, and C, and for all functions $f : A \to B$ and $g : B \to C$, if f and $g \circ f$ are one-to-one, then g is one-to-one.

Solution

1. There exist sets A, B, and C such that $A \cup C = B \cup C$ and $A \neq B$.
2. There exist sets A, B, and C and functions $f : A \to B$ and $g : B \to C$ such that f and $g \circ f$ are one-to-one, but g is not one-to-one. ■

Both parts of Example 1.5.13 are universal statements, so that their negations are existential statements. When one disproves a universal statement, we say that we are finding a *counterexample*.

Example 1.5.14 Show that the first statement in Example 1.5.13 is false.

Solution Let $A = \{1\}$, $B = \{2\}$, and $C = \{1, 2\}$. Then $A \cup C = B \cup C$, but $A \neq B$. ■

EXERCISE 1.5.15 Suppose each of the following statements could be disproved. By negating the statement, provide the statement that would be proved true. Assume that S is a given set.

(a) For all $x \in S$, $x \equiv x$.

(b) There exists a real number M such that $x \leq M$ for all $x \in S$.

(c) If $x_1 \parallel x_2$ and $x_2 \parallel x_3$, then $x_1 \parallel x_3$.

(d) For every function f and every set S, $f^{-1}[f(S)] \subseteq S$.

Properties of Real Numbers

It's time to begin applying the language and logic of Chapter 1 to proof writing. In this chapter, you will learn to write proofs of important properties of real numbers that follow from assumptions A1–A22 outlined in Chapter 0.

There are many principles of algebraic manipulation that you exploit without thinking. Here are some examples.

If $2x + 3 = 13$, then $2x = 10$.

If $2x = 10$, then $x = 5$.

If $a < b$ and $b < c$, then $a < c$.

If $|x| \leq 3$, then $-3 \leq x \leq 3$.

If $ab = 0$, then either $a = 0$ or $b = 0$.

Each of these properties can be proved as a theorem from assumptions A1–A22. Our tasks in this chapter are to understand how principles of algebraic manipulation follow from these assumptions and to learn to write proofs of such theorems.

2.1 Basic Algebraic Properties of Real Numbers

We begin by writing proofs of some of the most basic algebraic properties of real numbers. We will remind ourselves of the relevant assumptions as we need them. The properties of equality are of primary importance.

(A1) **Properties of equality:**

(a) For all $a \in \mathbb{R}$, $a = a$. (Reflexive)

(b) For all $a, b \in \mathbb{R}$, if $a = b$, then $b = a$. (Symmetric)

(c) For all $a, b, c \in \mathbb{R}$, if $a = b$ and $b = c$, then $a = c$. (Transitive)

The transitive property of equality allows us a convenience that we probably take for granted. When we write $a = b = c = d$, we interpret this as meaning that $a = d$. However, we must understand what is behind such a condensed statement. First, it includes the statement $(a = b) \wedge (b = c)$, from which we conclude $a = c$ by the transitive property. Second, it also includes the statement $c = d$, from which we conclude $a = d$, again by the transitive property.

2.1.1 Properties of Addition

There are six properties of real number addition that we take as axioms. Here they are again for reference.

(A2) **Addition is well defined:** For all $a, b, c, d \in \mathbb{R}$, if $a = b$ and $c = d$, then $a + c = b + d$.

(A3) **Closure property of addition:** For all $a, b \in \mathbb{R}, a + b \in \mathbb{R}$. The closure property also holds for $\mathbb{N}$, $\mathbb{W}$, $\mathbb{Z}$, and $\mathbb{Q}$.

(A4) **Associative property of addition:** For all $a, b, c \in \mathbb{R}$, $(a + b) + c = a + (b + c)$.

(A5) **Commutative property of addition:** For all $a, b \in \mathbb{R}, a + b = b + a$.

(A6) **Existence of an additive identity:** There exists a number $0 \in \mathbb{R}$ with the property that $a + 0 = a$ for all $a \in \mathbb{R}$.

(A7) **Existence of additive inverses:** For all $a \in \mathbb{R}$, there exists some $b \in \mathbb{R}$ such that $a + b = 0$. Such an element b is called an *additive inverse* of a and is typically denoted $-a$.

With these assumptions, we can now prove several properties of addition that follow from them. Read the proofs carefully, because you will mimic them to prove analogous properties for multiplication.

Theorem 2.1.1 (Cancellation of addition). For all $a, b, c \in \mathbb{R}$, if $a + c = b + c$, then $a = b$.

Proof. Pick $a, b, c \in \mathbb{R}$ and suppose $a + c = b + c$. By A7, there exists $-c \in \mathbb{R}$ such that $c + (-c) = 0$. Since addition is well defined (A2), we have that $[a + c] + (-c) = [b + c] + (-c)$, which by A4 we may write as $a + [c + (-c)] = b + [c + (-c)]$. This yields $a + 0 = b + 0$, which becomes $a = b$. □

We're going to present a second proof of Theorem 2.1.1, written in a somewhat different style. The previous proof reveals the thought process involved in constructing the proof, but it uses several disjoint equations. The next proof shows how you might clean up these equations to create one extended equation that is the result we want.

Proof 2. Suppose a, b, and c are real numbers such that $a + c = b + c$. By A7, there exists $-c \in \mathbb{R}$ such that $c + (-c) = 0$. Thus with several other properties from A1–A7 we have that

$$\begin{aligned} a &\overset{(A6)}{=} a + 0 \overset{(A2)}{=} a + [c + (-c)] \overset{(A4)}{=} (a + c) + (-c) \\ &\overset{(A2)}{=} (b + c) + (-c) \overset{(A4)}{=} b + [c + (-c)] \overset{(A2)}{=} b + 0 \overset{(A6)}{=} b \end{aligned} \tag{2.1}$$

By A1 (the transitive property of equality), $a = b$. □

Because addition is commutative, Theorem 2.1.1 also allows us to say that if $c + a = c + b$, then $a = b$. For if $c + a = c + b$, then

$$a + c = c + a = c + b = b + c$$

which by Theorem 2.1.1 implies $a = b$. Thus we have both *left cancellation* and *right cancellation*.

We will eventually get a little lazy about referencing the real number properties that we need to use, especially the three properties of equality. At first, however, we need to make sure for our own sakes that we understand precisely which ones we use and when we use them.

The next two results are uniqueness theorems. Recall what it means to say that a particular mathematical object is unique: If x_1 and x_2 are both assumed to have the desired property, then $x_1 = x_2$.

Theorem 2.1.2 The additive identity is unique.

Proof. Suppose 0_1 and 0_2 are both additive identities. Then

$$0_1 = 0_1 + 0_2 = 0_2$$

□

Assumption A7 guarantees that every real number a has *an* additive inverse $-a$ for which $a + -a = 0$. This is an axiom. But nothing about the axioms A1–A22 precludes (at first glance) that a real number might have more than one additive inverse. The next theorem shows that every real number has a unique additive inverse. Once we prove this, we will then be able to talk about *the* additive inverse of a. Understand what the uniqueness component of the existence of additive inverses would say. For all b_1, b_2, if b_1 and b_2 are both additive inverses of a, then $b_1 = b_2$.

Theorem 2.1.3 The additive inverse of a real number is unique.

Proof. Pick $a \in \mathbb{R}$, and suppose $b_1, b_2 \in \mathbb{R}$ are both additive inverses of a. Then $a + b_1 = 0$ and $a + b_2 = 0$. Thus by the transitive property of equality, $a + b_1 = a + b_2$, so that $b_1 = b_2$ by Theorem 2.1.1. Therefore the additive inverse of a is unique. □

Theorem 2.1.4 For every $a \in \mathbb{R}$, $-(-a) = a$.

Proof. Pick $a \in \mathbb{R}$. Then $-a$ exists in $\mathbb{R}$, where $a + (-a) = 0$. But then $-(-a)$ also exists in $\mathbb{R}$, where $(-a) + [-(-a)] = 0$. Thus $(-a) + [-(-a)] = a + (-a)$, which by cancellation yields $-(-a) = a$. □

Theorem 2.1.5 For all $a, b \in \mathbb{R}$, $-(a + b) = (-a) + (-b)$.

Proof. Pick $a, b \in \mathbb{R}$. Then $-a$ and $-b$ both exist in $\mathbb{R}$. By A3, $(-a) + (-b)$ is also a real number. Exploiting associativity and commutativity of addition, we have

$$(a + b) + [(-a) + (-b)] = [a + (-a)] + [b + (-b)] = 0 + 0 = 0 \tag{2.2}$$

Thus $-(a + b) = (-a) + (-b)$. □

One of the oldest gags around says that, if it quacks like a duck, it must be a duck. That is, the defining characteristic of ducks is that they quack. Now suppose you know that there is only one duck in the world, and you have just found an animal that quacks. Then, by golly, you have found the duck.

In Theorem 2.1.5, the duck we are looking for is $-(a + b)$. That is, we are looking for some real number, which, when added to $(a + b)$, yields zero. After all, that is the quacking feature of $-(a + b)$. Theorem 2.1.5 claims that the number $(-a) + (-b)$ quacks, and Eq. (2.2) is the demonstration of that claim. Finally, since the additive inverse of a real number is unique, we know that this quacking thing we have found is indeed $-(a + b)$.

Theorem 2.1.5 gives us the right to make a statement like the following:

The additive inverse of a sum is the sum of the additive inverses.

It is quite common in mathematics to investigate the truth of a statement of the form

$$\text{The } X \text{ of the } Y \text{ is equal to the } Y \text{ of the } X. \tag{2.3}$$

Statements of this form can be very powerful if true. However, sometimes such a statement is false. For example, let $X =$ "square root" and let $Y =$ "sum" and you have the statement $\sqrt{a + b} = \sqrt{a} + \sqrt{b}$.

Exercise 2.1.6 Calling on your experience from previous math courses, construct several statements with the same form as (2.3), some of which are true and some false.

The behavior of zero yields the next theorem.

Theorem 2.1.7 $-0 = 0$.

Proof. Since 0 is the additive identity, $0 + (-0) = -0$. Also, since -0 is the additive inverse of 0, by definition it satisfies $0 + (-0) = 0$. Thus we have $-0 = 0 + (-0) = 0$. □

2.1.2 Properties of Multiplication

Now it's time for you to write your first proofs. We begin with some theorems that are exactly analogous to Theorems 2.1.1–2.1.7, except that they are about multiplication instead of addition. We have to be a bit more careful with multiplication because multiplicative inverses are assumed to exist only for nonzero real numbers. Thus there are a few more results to prove for multiplication than there were for addition. First, we recall the relevant assumptions from Chapter 0. Note that we include the distributive property here (A14), as well as that very odd assumption that $1 \neq 0$ (A15).

(A8) **Multiplication is well defined:** For all $a, b, c, d \in \mathbb{R}$, if $a = b$ and $c = d$, then $ac = bd$.

(A9) **Closure property of multiplication:** For all $a, b \in \mathbb{R}$, $a \cdot b \in \mathbb{R}$. The closure property of multiplication also holds for $\mathbb{N}$, $\mathbb{W}$, $\mathbb{Z}$, and $\mathbb{Q}$.

(A10) **Associative property of multiplication:** For all $a, b, c \in \mathbb{R}$, $(ab)c = a(bc)$.

(A11) **Commutative property of multiplication:** For all $a, b \in \mathbb{R}$, $ab = ba$.

(A12) **Existence of a multiplicative identity:** There exists a number $1 \in \mathbb{R}$ with the property that $a \cdot 1 = a$ for all $a \in \mathbb{R}$.

(A13) **Existence of multiplicative inverses:** For all nonzero $a \in \mathbb{R}$, there exists some $b \in \mathbb{R}$ such that $ab = 1$. Such an element b is called a *multiplicative inverse* of a and is typically denoted a^{-1}.

(A14) **Distributive property of multiplication over addition:** For all $a, b, c \in \mathbb{R}$, $a(b + c) = ab + ac$.

(A15) $1 \neq 0$.

EXERCISE 2.1.8 For all $a, b, c \in \mathbb{R}$, if $ac = bc$ and $c \neq 0$, then $a = b$.

EXERCISE 2.1.9 The multiplicative identity is unique.

EXERCISE 2.1.10 The multiplicative inverse of a nonzero real number is unique.

Cancellation of addition comes in handy in the next theorem.

Theorem 2.1.11 For every real number a, $a \cdot 0 = 0$.

Proof. Pick $a \in \mathbb{R}$. By A6 and A14, we have that

$$0 + a \cdot 0 = a \cdot 0 = a \cdot (0 + 0) = a \cdot 0 + a \cdot 0 \tag{2.4}$$

By Theorem 2.1.1, we may cancel $a \cdot 0$ from both sides of Eq. (2.4) to have $a \cdot 0 = 0$. □

Occasionally, a theorem and its proof yield a special result that follows immediately but is important in its own right. Such a result is called a *corollary*. Sometimes the truth of the corollary is so transparent that it does not require its own proof. Sometimes, however, a word of explanation is in order.

Corollary 2.1.12 For all $a \in \mathbb{R}$, if $a \neq 0$, then $a^{-1} \neq 0$.

Proof. Pick any nonzero $a \in \mathbb{R}$. Since $a \cdot 0 = 0$ and $0 \neq 1$, then $a \cdot 0 \neq 1$. Thus $a^{-1} \neq 0$. □

Exercise 2.1.13 For all $a \neq 0$, $(a^{-1})^{-1} = a$.

Exercise 2.1.14 For all nonzero $a, b \in \mathbb{R}$, $(ab)^{-1} = a^{-1}b^{-1}$.[1]

Exercise 2.1.15 $1^{-1} = 1$.

Exercise 2.1.16 For all $a \in \mathbb{R}$, $(-1)a = -a$.[2]

Exercise 2.1.17 For all $a, b \in \mathbb{R}$, the following hold:

(a) $(-a)b = -(ab)$[3]
(b) $(-a)(-b) = ab$[4,5]

Suppose we have chosen some nonzero real number a. First take its additive inverse $-a$, and then take the multiplicative inverse of that to have $(-a)^{-1}$. Now start with a again, and perform the same two processes in reverse order. First take the multiplicative inverse of a to have a^{-1}; then take the additive inverse of that to have $-(a^{-1})$. You probably expect that these two processes, when done in reverse order, produce the same result, but if so, it must be demonstrated. Look closely at the proof of Theorem 2.1.18, for this kind of stunt can come in handy sometimes.

Theorem 2.1.18 For all nonzero real numbers a, $(-a)^{-1} = -(a^{-1})$.

Proof. Suppose $a \neq 0$. Then there exists $a^{-1} \in \mathbb{R}$, and by Exercise 2.1.17(b),

$$1 = a \cdot a^{-1} = (-a)[-(a^{-1})] \tag{2.5}$$

By Eq. (2.5), we see that the multiplicative inverse of $-a$ is $-(a^{-1})$. That is, $(-a)^{-1} = -(a^{-1})$. □

[1] Think about ducks.

[2] Ducks.

[3] You don't have to think about ducks. Just use previous results to transform the left-hand side into the right-hand side in several steps.

[4] Do you find yourself in need of the statement $(-1)(-1) = 1$? How do you know this is true? Look at the next hint if you have to.

[5] It's true because $(-1)(-1) = -(-1) = 1$ (Exercise 2.1.16 and Theorem 2.1.4).

EXERCISE 2.1.19 [Principle of zero products] If $ab = 0$, then either $a = 0$ or $b = 0$.[6]

EXERCISE 2.1.20 For all $a, b, c, d \in \mathbb{R}$, $(a + b)(c + d) = ac + ad + bc + bd$.

EXERCISE 2.1.21 Suppose we replace assumption A15 with the assumption that $1 = 0$. Show that, with this assumption, there are no nonzero real numbers.

EXERCISE 2.1.22 In this section two results in addition to Theorem 2.1.5 can be worded as "The X of the Y is equal to the Y of the X." Find them, and state them in this form.

2.2 Ordering Properties of the Real Numbers

In Chapter 0, we stated three assumptions about how the real numbers are ordered.

(A16) **Trichotomy law:** For any $a \in \mathbb{R}$, exactly one of the following is true:

(a) $a > 0$, in which case we say a is *positive*

(b) $a = 0$

(c) $0 > a$, in which case we say a is *negative*

(A17) For all $a, b \in \mathbb{R}$, if $a > 0$ and $b > 0$, then $a + b > 0$.

(A18) For all $a, b \in \mathbb{R}$, if $a > 0$ and $b > 0$, then $ab > 0$.

We address the *sign* of real numbers by assuming that every nonzero real number can, by the trichotomy law, be compared to zero in one way or another by the symbol $>$. If $a > 0$, we say that a is *positive*, and if $0 > a$ (or $a < 0$), we say that a is *negative*. Notationally, we write the positive and negative real numbers as $\mathbb{R}^+$ and $\mathbb{R}^-$, respectively. We have therefore split the real numbers into three pieces, $\mathbb{R}^+$, $\{0\}$, and $\mathbb{R}^-$, and we have two closure assumptions about $\mathbb{R}^+$.

We can now assign meaning to the statement $a > b$, as we did in Definition 0.2.1. We restate that definition here for reference. Recall from Chapter 0 that $a - b = a + (-b)$ by definition.

Definition 2.2.1 Given real numbers a and b, we say that $a > b$ provided $a - b > 0$. The statement $a < b$ means $b > a$. The statement $a \geq b$ means that either $a > b$ or $a = b$. Similarly, $a \leq b$ means either $a < b$ or $a = b$.

So anytime we see a statement $a > b$, it means precisely that $a - b > 0$. Similarly, if someone asks you to demonstrate that $a > b$, we can do it by showing $a - b > 0$. Definition 2.2.1, with Theorem 2.1.4, allows us to prove the following.

[6] See Exercise 1.2.18(h).

Theorem 2.2.2 If $a > b$, then $-a < -b$.

Proof. Suppose $a > b$. Then $a - b > 0$. But $a = -(-a)$, so that $-(-a) - b > 0$, which we may write as $-(-a) + (-b) > 0$. By commutativity, $(-b) + [-(-a)] > 0$, or $-b - (-a) > 0$, so that $-b > -a$. Hence $-a < -b$. □

Corollary 2.2.3 Suppose c is any real number. Then $c > 0$ if and only if $-c < 0$.

Proof. ($\Rightarrow$) Suppose $c > 0$. Then we may let $a = c$ and $b = 0$ in Theorem 2.2.2 to have $-c < -0 = 0$.

($\Leftarrow$) Suppose $-c < 0$. Then we may let $a = 0$ and $b = -c$ in Theorem 2.2.2 to have $-0 < -(-c)$, or $c > 0$. □

Assumption A15 states that $1 \neq 0$, which might seem silly, but actually is an essential assumption. In Exercise 2.1.21, you showed that the assumption $1 = 0$ causes the entire set of real numbers to collapse to the single element set $\{0\}$. Since $1 \neq 0$ by assumption, the trichotomy law implies that either $1 > 0$ or $1 < 0$. To tackle the next exercise, try a proof by contradiction.

EXERCISE 2.2.4 $1 > 0$.

EXERCISE 2.2.5 Prove the following.

(a) If $a < b$, then $a + c < b + c$.

(b) If $a < b$ and $c = d$, then $a + c < b + d$.

(c) If $a < b$ and $c < d$, then $a + c < b + d$.

(d) If $a \leq b$ and $c \leq d$, then $a + c \leq b + d$.[7]

(e) If $a < b$ and $b < c$, then $a < c$.

(f) If $a > b$ and $b > c$, then $a > c$.

(g) If $a < 0$ and $b < 0$, then $a + b < 0$.

(h) If $a > 0$ and $b < 0$, then $ab < 0$.

(i) If $a < 0$ and $b < 0$, then $ab > 0$.

(j) If $a < b$ and $c > 0$, then $ac < bc$.

(k) If $a < b$ and $c < 0$, then $ac > bc$.

EXERCISE 2.2.6 For a real number a, write $a^2 = a \cdot a$. Prove the following.

(a) If $0 < a < b$, then $a^2 < b^2$.

(b) If $a < b < 0$, then $a^2 > b^2$.

[7] Case it out.

(c) For all real numbers a, $a^2 \geq 0$.

(d) The equation $x^2 = -1$ has no real number solution x.

Exercise 2.2.7 Prove the following.

(a) 0^{-1} does not exist in $\mathbb{R}$.

(b) If $a > 0$, then $a^{-1} > 0$.

(c) If $a < 0$, then $a^{-1} < 0$.

(d) $c > 1$ if and only if $0 < c^{-1} < 1$.[8]

(e) If $a > 0$ and $c > 1$, then $a/c < a$.[9]

(f) If a and b are integers such that $ab = 1$, then $a = b = \pm 1$.[10]

Exercise 2.2.8 If a and b are nonzero real numbers and $a < b$, does it follow that $1/a > 1/b$? Use results from Exercises 2.2.5 and 2.2.7 to state and prove the relationship between $1/a$ and $1/b$ depending on the signs of a and b.

Exercise 2.2.9 Prove that if $a < b$, then $a < (a + b)/2 < b$. (How do you know that $2 > 0$? What is 2 anyway?)

2.3 Absolute Value

You are certainly familiar with the *absolute value* of real numbers:

Definition 2.3.1 For a real number x, we define $|x|$, the *absolute value* of x, by

$$|x| = \begin{cases} x, & \text{if } x > 0 \\ 0, & \text{if } x = 0 \\ -x, & \text{if } x < 0 \end{cases} \tag{2.6}$$

The fact that $|x|$ is defined in three cases means that proofs of theorems often have to address the cases separately. In many situations, two cases suffice, by noting that the following is equivalent to Definition 2.3.1.

$$|x| = \begin{cases} x, & \text{if } x \geq 0 \\ -x, & \text{if } x < 0 \end{cases} \tag{2.7}$$

[8] To show $0 < c^{-1} < 1$, you must show the two inequalities $0 < c^{-1}$ and $c^{-1} < 1$ separately.

[9] Recall $a/c = ac^{-1}$.

[10] There are five cases: $a = -1$, $a = 0$, $a = 1$, $a > 1$, and $a < -1$. Three are impossible. Remember that 1 is the smallest positive integer.

Absolute value is a very important measure of the *size* of a real number. We can make two observations right off the bat.

(N1) For every real number x, $|x| \geq 0$.

(N2) $|x| = 0$ if and only if $x = 0$.

First let's look at property N1. Pick any real number x. If $x \geq 0$, then $|x| = x \geq 0$. If $x < 0$, then $|x| = -x > 0$ by Corollary 2.2.3. Thus in any case, $|x| \geq 0$.

Property N2 follows directly from Corollary 2.2.3 and the trichotomy law. The reason we point these out is that properties N1–N2 are two of the three defining properties of a *norm*, a very important term in analysis. Given a set (perhaps of numbers, functions, sets) a norm is a measure of the size of its elements. There is a third property of a norm that absolute value has, and we will see it in Exercise 2.3.8.

First let's explore some of the simplest and most familiar behaviors of absolute value. We will prove a few, either wholly or in part, and leave some to you as exercises. The first one is really easy, so it's all yours. You'll probably want to use the three cases in Eq. (2.6).

EXERCISE 2.3.2 For all real numbers x, $|-x| = |x|$.

Theorem 2.3.3 Let $a \geq 0$ be a given real number. Then $|x| = a$ if and only if $x = \pm a$.

Proof. Let $a \geq 0$ be given.

($\Rightarrow$) Suppose $|x| = a$. If $x \geq 0$, then $x = |x| = a$. If $x < 0$, then $x = -|x| = -a$. In either case, $x = \pm a$.

($\Leftarrow$) Suppose $x = \pm a$. First note that if $a = 0$, then $|x| = x = a$. So suppose $a > 0$. For the case $x = a$, we have that $x > 0$, so that $|x| = x = a$. For the case $x = -a$, we have that $x < 0$. Thus $|x| = -x = a$. In all these cases, $|x| = a$. □

The next result is an if-and-only-if theorem, where each direction can be handled by considering the two cases $x \geq 0$ and $x < 0$.

Theorem 2.3.4 Let $a > 0$ be a given real number. Then $|x| < a$ if and only if $-a < x < a$.

Proof. Let $a > 0$ be given.

($\Rightarrow$) Suppose $|x| < a$. If $x \geq 0$, then

$$-a < 0 \leq |x| = x < a \tag{2.8}$$

On the other hand, if $x < 0$, we may write $-|x| > -a$ to have

$$-a < -|x| = -(-x) = x < 0 < a \tag{2.9}$$

In either case, we have $-a < x < a$.

($\Leftarrow$) Now suppose $-a < x < a$. If $x \geq 0$, then $|x| = x < a$. If $x < 0$, we note that $-a < x$ is equivalent to $-x < a$, and we have $|x| = -x < a$. In either case, $|x| < a$.

□

We may apply Exercise 1.2.18(j) to derive a corollary to Theorems 2.3.3 and 2.3.4. Define the following statements:

$$p: \ |x| = a \quad q: \ x = \pm a \quad r: \ |x| < a \quad s: \ -a < x < a \tag{2.10}$$

By Exercise 1.2.18(j), $(p \leftrightarrow q) \wedge (r \leftrightarrow s)$ is stronger than $(p \vee r) \leftrightarrow (q \vee s)$. We can therefore combine Theorems 2.3.3 and 2.3.4 to have the following.

Corollary 2.3.5 Let $a > 0$ be a given real number. Then $|x| \leq a$ if and only if $-a \leq x \leq a$.

You can prove the next result by casing it out, if you like. Or you can apply Corollary 2.3.5 and some sneaky logic to prove it quickly.

Exercise 2.3.6 Let $a > 0$ be a given real number. Then $|x| > a$ if and only if either $x > a$ or $x < -a$.

Exercise 2.3.7 For all real numbers x, $-|x| \leq x \leq |x|$.

Now let's state the third property of a norm, and show that $|x|$ has this important property. Exercise 2.2.5(d), Corollary 2.3.5, and Exercise 2.3.7 will provide just the right mathematical machinery to prove the following.

Exercise 2.3.8 [N3: Triangle Inequality] For all real numbers x and y,

$$|x + y| \leq |x| + |y|.^{11} \tag{2.11}$$

Another triangle-type inequality can be proved in two lines from Exercise 2.3.8, if you can see how to apply it creatively.

Exercise 2.3.9 For all real numbers x and y, $|x - y| \geq |x| - |y|$.[12,13]

There is one more triangle-type inequality that is important in analysis. The proof requires some sneaky application of several of the results we have shown so far.

[11] Find M such that $-M \leq x + y \leq M$ and apply Corollary 2.3.5.

[12] One of the mathematician's most useful tricks is knowing when and how to add zero. Start with Exercise 2.3.8, and reassign the roles of x and y. Don't look at the next hint unless you have to.

[13] Start with $|x|$ by itself and add zero inside the absolute value.

EXERCISE 2.3.10 For all real numbers x and y, $\big||x| - |y|\big| \leq |x - y|$.[14]

Exercises 2.3.8–2.3.10 reveal how addition and subtraction relate to absolute value. The final exercise for this section reveals how multiplication and division relate to absolute value.

EXERCISE 2.3.11 Suppose x and y are real numbers. Then the following are true.

(a) $|xy| = |x|\,|y|$.

(b) If $y \neq 0$, then $\left|y^{-1}\right| = |y|^{-1}$.[15]

(c) If $y \neq 0$, then $|x/y| = |x|/|y|$.

2.4 The Division Algorithm

This section is devoted to one very important theorem about the integers. Even though it is called the *division algorithm*, it is not an algorithm in the sense of computer science, where we think of an algorithm as a step-by-step process for performing a task. Instead, the division algorithm is merely a unique existence theorem, but a very useful one.

One way to assign meaning to the words *even integer* and *odd integer* is the following.

Definition 2.4.1 If n is an integer, we say that n is *even* provided there exists an integer k such that $n = 2k$. We say that n is *odd* provided there exists an integer k such that $n = 2k + 1$.

Let's prove some easy results about even and odd integers. Notice how the proof of the first one appeals to the definition of even integer in both directions. First, when we are given that a particular integer n is even, then we can claim the existence of some integer k_1 such that $n = 2k_1$. Then, when we need to demonstrate that some integer p is even, we must be able to find some integer k_2 such that $p = 2k_2$.

We present two slightly different proofs of the next theorem to illustrate subtly different tastes for the flow of a proof.

Theorem 2.4.2 Let m and n be integers, at least one of which is even. Then mn is even.

Proof 1. Pick integers m and n, and suppose (without loss of generality) that n is even. Then there exists integer k_1 such that $n = 2k_1$. Thus we have that

[14] Apply Corollary 2.3.5 by writing the expression in Exercise 2.3.9 in two ways, once switching the roles of x and y.

[15] Use a technique like the proof of Theorem 2.1.18.

$mn = m(2k_1) = 2(mk_1)$, so that we may let $k_2 = mk_1$ (which is an integer) to have $mn = 2k_2$. Thus mn is even. □

Proof 2. Pick integers m and n, and suppose (without loss of generality) that n is even. Then there exists integer k_1 such that $n = 2k_1$. Let $k_2 = mk_1$, which by closure is also an integer. Then

$$mn = m(2k_1) = 2(mk_1) = 2k_2 \tag{2.12}$$

Thus mn is even. □

Corollary 2.4.3 If n is an even integer, then n^2 is even.

Proof. Let $m = n$ in Theorem 2.4.2. □

EXERCISE 2.4.4 If m and n are odd integers, then mn is odd.

Corollary 2.4.5 If n is an odd integer, then n^2 is odd.

EXERCISE 2.4.6 Prove the following.

(a) The sum of two even integers is even.

(b) The sum of two odd integers is even.

(c) The sum of an even integer and an odd integer is odd.

At this point, you're likely thinking that Definition 2.4.1 addresses the only two possible situations that can happen in the integers. After all, isn't every integer either even or odd, and not both? Well, maybe, but how do you know that? How do you know that every integer can be written either in the form $2k$ or $2k + 1$, but not both? This, and a whole lot more, is addressed by the division algorithm.

Ever since elementary school, you have been familiar with the idea of dividing an integer b by another integer $a > 0$ to produce a *quotient* q and *remainder* r. One way you could express the results of your division calculation in an equation is to write

$$b = aq + r \tag{2.13}$$

where $0 \le r < a$. For example, if $a = 12, b = 88$, we can write $88 = (12)(7) + 4$, and if $a = 6, b = -13$, we have $-13 = (6)(-3) + 5$. One nice thing about the form of Eq. (2.13) is that every number involved is an integer. What resembles division of integer by integer is written as an integer equation without having to resort to rational numbers. The division algorithm says that the existence of a quotient q and a remainder $0 \le r < a$ are guaranteed. Even more, they are unique. The theorem is a surprisingly useful one, as you will come to see. In its proof we need an assumption from Chapter 0 that we have not used yet: the well-ordering principle (WOP).

(A19) **Well-ordering principle:** Any non-empty subset of whole numbers has a smallest element. That is, if A is non-empty, then there is some number $a \in A$ with the property that $a \leq x$ for all $x \in A$.

Theorem 2.4.7 (Division Algorithm). Let a and b be integers, where $a > 0$. Then there exist unique integers q and r such that $b = aq + r$ and $0 \leq r < a$.

We provide the existence part of the proof here. Then you will show uniqueness as an exercise. In showing existence for the case $b \geq 0$, we use the set

$$S = \{b - aq : q \in \mathbb{Z},\ b - aq \geq 0\} \tag{2.14}$$

which is merely the set of all whole numbers you can generate by subtracting integer multiples of a from b. For example, if $a = 12$ and $b = 105$, then

$$S = \{\dots, 129, 117, 105, 93, 81, 69, 57, 45, 33, 21, 9\} \tag{2.15}$$

The last element of S is $105 - 12q$, where $q = 8$. Notice that it is the smallest element of S and is strictly less than 12. By constructing the corresponding S according to Eq. (2.14) for an arbitrary $a > 0$ and $b \geq 0$, we can apply the WOP to S, then show that this smallest element is the value of r we're after. Here is the proof of the existence part of Theorem 2.4.7, in as much detail as we promised.

Proof. (Existence) Let a and b be integers, where $a > 0$. First, we consider the case $b \geq 0$. Define the set S as in Eq. (2.14). By definition, $S \subseteq \mathbb{W}$, and since $b \geq 0$, we may let $q = 0$ to see that $b \in S$, so that S is non-empty. By the WOP, S has a smallest element, which we may denote r, and note that r is of the form $b - aq$ for some integer q. Thus we have that $b = aq + r$, where $r \geq 0$, and we must show that $r < a$. To show $r < a$, suppose $r \geq a$ (from which we will arrive at a contradiction). Then $b - a(q + 1) = b - aq - a = r - a \geq 0$, so that $b - a(q + 1) \in S$. Also, $b - a(q + 1) = b - aq - a < b - aq = r$. Thus $b - a(q + 1)$ is an element of S that is smaller than r. This is a contradiction, because r was chosen to be the smallest element of S. Thus it must be that $r < a$.

Now suppose $b < 0$, so that $-b > 0$. Applying the previous case to $-b$, there exist integers q_1 and r_1 such that $(-b) = aq_1 + r_1$ and $0 \leq r_1 < a$. Now if $r_1 = 0$, then $-b = aq_1$, so that $b = a(-q_1)$. Thus we may let $q = -q_1$ and have that $b = aq + 0$. So suppose that $0 < r_1 < a$. Then $b = a(-q_1) - r_1 = a(q_1 + 1) - (a - r_1)$. Let $q = q_1 + 1$ and $r = a - r_1$. Then q and r are integers, and the fact that $0 < r_1 < a$ implies that $-a < -r_1 < 0$, so that $0 < a - r_1 < a$. Thus $0 < r < a$. □

Exercise 2.4.8 Show uniqueness of q and r in Theorem 2.4.7.[16]

[16] Suppose $q_1, q_2, r_1, r_2 \in \mathbb{Z}$ are such that $b = aq_1 + r_1$ $(0 \leq r_1 < a)$ and $b = aq_2 + r_2$ $(0 \leq r_2 < a)$. There are two useful things you can say about $r_2 - r_1$.

Exercise 2.4.9 Appeal to the division algorithm to explain precisely why every integer is either even or odd, but not both.

2.5 Divisibility and Prime Numbers

Suppose a and b are nonzero integers and we apply the division algorithm to write $b = aq + r$. The situation where $r = 0$ motivates a new term.

Definition 2.5.1 Suppose a and b are nonzero integers. We say that *a divides b*, written $a \mid b$, provided there exists an integer k such that $b = ak$. We also say that a is a *divisor* of b. If $a \mid b$ where $a \notin \{\pm 1, \pm b\}$, we call a a *proper* divisor of b. If a does not divide b, we write $a \nmid b$.

Notice that if a does not divide b, then writing $b = aq + r$ according to the division algorithm will imply $0 < r < a$. First, here are some really easy results about divisibility for you to prove.

Exercise 2.5.2 If a is a nonzero integer, then $a \mid a$.

Exercise 2.5.3 If a, b, and c are nonzero integers such that $a \mid b$ and $b \mid c$, then $a \mid c$.

An expression of the form $mb + nc$ is called a *linear combination* of b and c. The next exercise says that if a divides both b and c, then it divides any integer linear combination of them.

Exercise 2.5.4 If a, b, and c are nonzero integers such that $a \mid b$ and $a \mid c$, then $a \mid (mb + nc)$ for all integers m and n.

Exercise 2.5.5 If a and b are nonzero integers such that $a \mid b$ and $b \mid a$, then $a = \pm b$.[17]

Some important and useful results in algebra stem from what we call the *greatest common divisor* (gcd) of two nonzero integers a and b. First we must define what we mean by the gcd of two nonzero integers. We do this by stating two criteria that a gcd must satisfy, one motivated by the word *common*, and the other motivated by the word *greatest*. Then we show that such a thing exists uniquely and can be written in a somewhat surprising way. So that we will have uniqueness of $\gcd(a, b)$, we insist that any integer that hopes to qualify as gcd must be positive.

Definition 2.5.6 Suppose a and b are nonzero integers, and suppose g is a positive integer with the following properties:

[17] See Exercise 2.2.7(f).

(D1) $g \mid a$ and $g \mid b$

(D2) If h is any positive integer such that $h \mid a$ and $h \mid b$, then $h \mid g$ also.

Then g is called a *greatest common divisor* of a and b, and is denoted $\gcd(a, b)$.

Some remarks about Definition 2.5.6 are in order. First, nothing about the definition of $\gcd(a, b)$ (or of any term for that matter) guarantees that any such thing exists. It merely lays out criteria by which some positive integer earns the right to be called $\gcd(a, b)$. Second, property D1 states that any integer that might qualify as $\gcd(a, b)$ will in fact be a common divisor of a and b. Third, we need to explain our choice of property D2 as the other criterion, for it might seem a bit unnatural. If property D2 is supposed to describe what it means for g to be greatest of all the common divisors of a and b, you might think a more natural way to say it would be

(D2′) If h is any positive integer such that $h \mid a$ and $h \mid b$, then $h \leq g$.

Well, we could do it that way, but we don't. Here's why. It all centers around how you decide to measure greatness. In the integers, we can measure relative greatness by $\leq$. But as you will see in your later work in algebra, there are other more abstract settings where we discuss divisibility, then ask about a gcd, but we do not necessarily have a notion of $\leq$ to use as our criterion for greatness. The point is that it is better to develop a criterion for greatness of a divisor that stays within the context of divisibility rather than relying on some external measure like $\leq$, which might or might not already exist. All this is to say that property D2 is our measure of greatness among all common divisors of a and b. If g is going to qualify as $\gcd(a, b)$, it has the property that any positive integer h that comes down the pike also having property D1 cannot, in some sense, be greater than g. Our way of saying h is no greater than g is that $h \mid g$.

In the past, you have probably found $\gcd(a, b)$ by breaking a and b down into their prime factorizations to see how many 2s, 3s, 5s, and so on, that you can factor out of each. Then you built the gcd to contain just the right number of each prime factor represented. This might work well practically, but (1) we have no logical basis for this in the work we have done so far because we have not discussed prime numbers yet, and (2) it is not particularly useful as leverage in later theorems.

So, given nonzero integers a and b, how do we even know that there exists some positive integer having properties D1–D2? And if such an integer does exist, how many can there be? The following theorem claims that $\gcd(a, b)$ exists uniquely. Furthermore, hidden inside its proof is some additional information about $\gcd(a, b)$ that comes in handy later. Most of the verification of the details is left to you as an exercise.

Theorem 2.5.7 Suppose a and b are nonzero integers. Then there exists a unique positive integer g with the following properties:

(D1) $g \mid a$ and $g \mid b$.

(D2) If h is any positive integer such that $h \mid a$ and $h \mid b$, then $h \mid g$ also.

Proof. Let a and b be nonzero integers, and define

$$S = \{ma + nb : m, n \in \mathbb{Z}, ma + nb > 0\} \tag{2.16}$$

That is, S is the set of all *positive* integer linear combinations of a and b. First, S is not empty, for depending on the signs of a and b, we may let $m, n = \pm 1$ to produce some positive value of $ma + nb$. By the WOP, S contains a smallest element g, which may be written in the form $g = m_0 a + n_0 b$ for some integers m_0 and n_0. By Exercise 2.5.8, g has properties D1–D2, and if g_1 and g_2 are both positive integers that have properties D1–D2, then $g_1 = g_2$. □

EXERCISE 2.5.8 Let $g = m_0 a + n_0 b$ be the smallest element of S as defined in Eq. (2.16). Show the following to complete the proof of Theorem 2.5.7.

(a) $g \mid a$. (The proof that $g \mid b$ is identical.)[18, 19]
(b) If h is any positive integer with the properties that $h \mid a$ and $h \mid b$, then it must be that $h \mid g$.
(c) If g_1 and g_2 are positive integers with properties D1–D2, then $g_1 = g_2$.[20]

The proof of Theorem 2.5.7 yields a serendipity about $\gcd(a, b)$, in that it can be written as an integer linear combination of a and b, and the smallest such positive expression is in fact the gcd. This can be particularly helpful if $\gcd(a, b) = 1$, in which case we say that a and b are *relatively prime*. Immediately, we can see that if a and b are relatively prime, then there exist integers m and n such that $ma + nb = 1$. Furthermore, if it is possible to find an integer linear combination of a and b that equals 1, then clearly this linear combination is the smallest such positive linear combination. Thus a and b are relatively prime.

EXERCISE 2.5.9 Show that 14 and 33 are relatively prime by determining integers m and n such that $14m + 33n = 1$.

Now we are ready to define prime numbers and investigate a few of their important properties. A positive integer p is defined to be prime if it has precisely two *distinct* positive integer divisors. If this is true, then these divisors must be 1 and p. Note that this definition excludes 1 from being prime.

Theorem 2.5.10 Suppose a is a nonzero integer and p is prime. Then either a and p are relatively prime or $p \mid a$.

[18] Apply the division algorithm to a and g and show that $r > 0$ is impossible.

[19] To suppose $r > 0$ leads to a contradiction because it produces an element of S that is smaller than g.

[20] Theorem 2.5.5 should come in handy.

Proof. Let a be a nonzero integer, and let p be prime. Since the only positive integer divisors of p are 1 and p, then either $\gcd(a, p) = 1$ or $\gcd(a, p) = p$. In the former case, a and p are relatively prime. In the latter case, p satisfies property D1, so that $p \mid a$. □

EXERCISE 2.5.11 Suppose a and b are nonzero integers, p is prime, and $p \mid ab$. Then either $p \mid a$ or $p \mid b$.[21]

As an immediate corollary, by letting $a = b$ in Exercise 2.5.11, we have the following.

Corollary 2.5.12 If $p \mid a^2$, then $p \mid a$.

EXERCISE 2.5.13 Let $a_1, a_2, \ldots, a_n$ be nonzero integers.

(a) Suppose a is an integer such that $a \mid a_j$ for all $1 \leq j \leq n$. Show that a divides any integer linear combination of $a_1, \ldots, a_n$.
(b) Motivated by Definition 2.5.6 create a definition of $\gcd(a_1, a_2, \ldots, a_n)$.
(c) State and prove a parallel to Theorem 2.5.7.

EXERCISE 2.5.14 If $n \geq 3$ is an odd integer, then $8 \mid (n^2 - 1)$.[22]

EXERCISE 2.5.15 If n is an integer such that $n \geq 2$ and $3 \mid (n - 1)$, then $3 \mid (n^2 - 1)$.

[21] See Exercise 1.2.18(h).

[22] If k is an integer, then either k or $k + 1$ is even.

Sets and Their Properties

All of the results in Chapter 2 were about familiar sets of numbers, in particular their algebraic properties arising from the binary operations of addition and multiplication. This chapter will also address properties of real numbers, but within a broadened context of their familiar subsets ($\mathbb{N}$, $\mathbb{W}$, $\mathbb{Z}$, and $\mathbb{Q}$) and sets in general.

3.1 Set Terminology

First, we return to the set terms from Chapter 0 and introduce several others. Given a set A in the context of a universal set U, we often illustrate with a *Venn diagram*, as in Figure 3.1. Shading a Venn diagram can sometimes be helpful to illustrate sets under discussion.

Definition 3.1.1 Given a set A, we define the *complement* of A, denoted A^C, in the following way.

$$A^C = \{x : x \in U \text{ and } x \notin A\} \tag{3.1}$$

Thus $x \in A^C$ if and only if $x \notin A$. See Figure 3.2.

Given two sets A and B, a general Venn diagram would look like that in Figure 3.3. We can now define several important binary operations to create new sets from A and B.

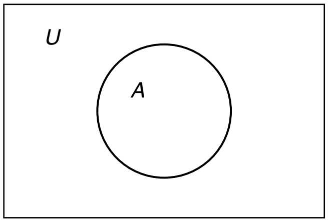

Figure 3.1 Basic Venn diagram.

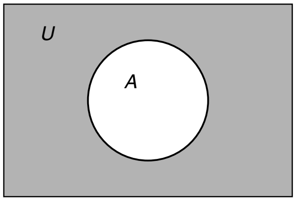

Figure 3.2 Venn diagram illustrating A^C.

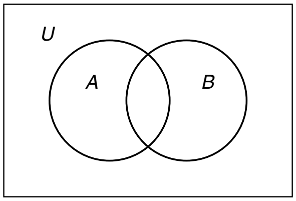

Figure 3.3 Generic Venn diagram with two sets, A and B.

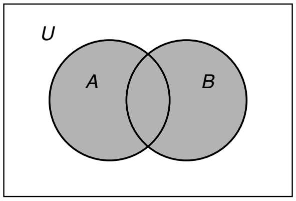

Figure 3.4 Venn diagram illustrating $A \cup B$.

Definition 3.1.2 Given two sets A and B, we define their *union* $A \cup B$ in the following way.

$$A \cup B = \{x : x \in A \text{ or } x \in B\}$$

Thus $x \in A \cup B$ if and only if $x \in A$ or $x \in B$. See Figure 3.4.

Definition 3.1.3 Given two sets A and B, we define their *intersection* $A \cap B$ in the following way.

$$A \cap B = \{x : x \in A \text{ and } x \in B\}$$

Thus $x \in A \cap B$ if and only if $x \in A$ and $x \in B$. See Figure 3.5.

Definition 3.1.4 Given two sets A and B, we define their *difference* $A - B$ in the following way.

$$A - B = A \cap B^C = \{x : x \in A \text{ and } x \notin B\}$$

Thus $x \in A - B$ if and only if $x \in A$ and $x \notin B$. See Figure 3.6.

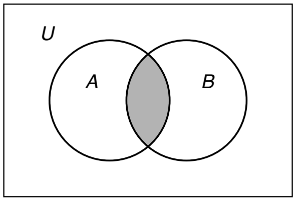

Figure 3.5 Venn diagram illustrating $A \cap B$.

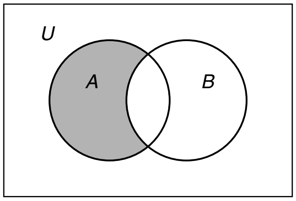

Figure 3.6 Venn diagram illustrating $A - B$.

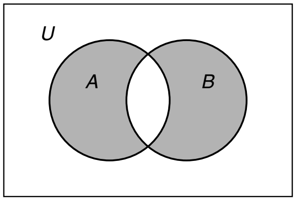

Figure 3.7 Venn diagram illustrating $A \Delta B$.

Definition 3.1.5 Given two sets A and B, we define the *symmetric difference* of A and B in the following way.

$$A \Delta B = (A \cap B^C) \cup (A^C \cap B) = \{x : x \in A \cap B^C \text{ or } x \in A^C \cap B\} \tag{3.2}$$

Thus $x \in A \Delta B$ if and only if $x \in A - B$ or $x \in B - A$. See Figure 3.7. Note that $x \in A \Delta B$ if x is in precisely one of sets A and B.

EXERCISE 3.1.6 Construct the negation of the following statements.

(a) $x \in A^C$

(b) $x \in A \cup B$

(c) $x \in A \cap B$

(d) $x \in A - B$

(e) $x \in A \Delta B$

Now we define some terms of comparison for sets.

Definition 3.1.7 Suppose A and B are sets. We say that A is a *subset* of B, written $A \subseteq B$, provided that for all x, if $x \in A$, then $x \in B$. That is,

$$(A \subseteq B) \Leftrightarrow (\forall x)(x \in A \rightarrow x \in B) \Leftrightarrow (\forall x \in A)(x \in B) \tag{3.3}$$

Exercise 3.1.8 In the spirit of Definition 3.1.7, construct a definition of $A \subset B$ from its definition on p. 2.

The relationship $A \subseteq B$ can be displayed in the Venn diagram in Figure 3.8. This general diagram is not intended to imply that the region outside A but inside B contains any elements.

Definition 3.1.9 If A and B are sets, we say that $A = B$ provided $A \subseteq B$ and $B \subseteq A$. That is,

$$\begin{aligned} A = B &\Leftrightarrow (A \subseteq B) \wedge (B \subseteq A) \\ &\Leftrightarrow (\forall x)(x \in A \rightarrow x \in B) \wedge (\forall x)(x \in B \rightarrow x \in A) \end{aligned} \tag{3.4}$$

Exercise 3.1.10 Construct a negation of the following statements.

(a) $A \subseteq B$

(b) $A \subset B$

(c) $A = B$

(d) $A \subseteq B \cup C$

(e) $A \cap B \subseteq C \cup D$

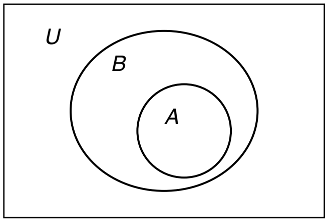

Figure 3.8 Venn diagram with $A \subseteq B$.

(f) For all sets A, B, and C, if $A \cup C \subseteq B \cup C$, then $A \subseteq B$.

(g) For all sets A and B, if $A \subseteq B$, then $A \cap B = A$.

Definition 3.1.11 Given two sets A and B, we say that they are *disjoint* provided $A \cap B$ is empty.

Exercise 3.1.12 If A and B are not disjoint, then there exists an element in $A \cap B$. By negating this, construct an if-then statement that is equivalent to the statement that A and B are disjoint.

Exercise 3.1.13 In this exercise, A and B represent sets, and x is an unspecified element. In each part of this exercise, you are presented with two statements, one of which is stronger than the other. State which one is stronger.

(a)	$x \in A$	$x \in A \cup B$
(b)	$x \in A$	$x \in A \cap B$
(c)	$x \in A$	$x \in A - B$
(d)	$x \in A \Delta B$	$x \in A - B$
(e)	$A \subseteq B$	$A = B$
(f)	$A \subset B$	$A \subseteq B$

3.2 Proving Basic Set Properties

Now we turn to proofs of basic set properties. At this point we fulfill the promise we made in Section 1.5 to provide an example of a theorem proved with a tautology. A theorem about set equality is a natural place to do this. We will set up the logical statements, then leave the truth table to you as an exercise. Then we will write a proof in paragraph form. Here is the theorem we will prove in these ways. Note that it is a universal statement about all sets A and B that satisfy $A \subseteq B$. We omit the universal quantifiers.

Theorem 3.2.1 If $A \subseteq B$, then $A \cap B = A$.

In this set theorem, a Venn diagram can convince you that the theorem is true. If the Venn diagram is drawn as in Figure 3.8, then the fact that $A \cap B = A$ is apparent. But Venn diagrams do not constitute a proof at this point in your mathematical career. You have a long way to go before you earn the right to say "Proof by picture."

Define the following logical phrases.

$$p: x \in A \tag{3.5}$$

$$q: x \in B \tag{3.6}$$

How do we state the hypothesis of Theorem 3.2.1 in terms of these statements p and q?[1] Also, how do we state the conclusion $A \cap B = A$ in the same terms?[2] The answers:

$$\overbrace{(p \to q)}^{\text{hypothesis}} \to \overbrace{[(p \wedge q) \to p] \wedge [p \to (p \wedge q)]}^{\text{conclusion}} \tag{3.7}$$

Exercise 3.2.2 Construct a truth table for (3.7) to verify it is a tautology.

Exercise 3.2.3 Using p and q as defined in (3.5) and (3.6), translate the following theorem into symbolic form and prove it with a truth table.

If $A \cup B = B$, then $A \subseteq B$.

Now we present a proof of Theorem 3.2.1 in mathematical prose. To prove two sets are equal, we must show that each is a subset of the other. Furthermore, to show that one set is a subset of another, we must show that every element of the former is also an element of the latter. Proofs of this sort are often called *element-chasing* proofs. They show two sets are equal by chasing an arbitrarily chosen element from one side to the other, then back.

Proof of Theorem 3.2.1. Suppose $A \subseteq B$. To show $A \cap B = A$, show that $A \cap B \subseteq A$ and $A \subseteq A \cap B$.

($A \cap B \subseteq A$): Pick $x \in A \cap B$. Then $x \in A$ and $x \in B$. Since $x \in A$, we have that $A \cap B \subseteq A$.

($A \subseteq A \cap B$): Pick $x \in A$. Then, since $A \subseteq B$, we have that $x \in B$ also. Therefore, since $x \in A$ and $x \in B$, it follows that $x \in A \cap B$. Thus $A \subseteq A \cap B$.

Since we have shown that $A \cap B \subseteq A$ and $A \subseteq A \cap B$, we have demonstrated that $A \cap B = A$. □

This proof of Theorem 3.2.1 is intentionally wordier than it needs to be, so that you can see how its structure derives from the definitions of the terms it uses. A better, more succinct version would read as follows.

Proof of Theorem 3.2.1 (Succinct). Suppose $A \subseteq B$.

($\subseteq$) Pick $x \in A \cap B$. Then $x \in A$, so that $A \cap B \subseteq A$.

($\supseteq$) Pick $x \in A$. Then since $A \subseteq B$, it is also true that $x \in B$, so that $x \in A \cap B$. Thus $A \subseteq A \cap B$.

Since $A \cap B \subseteq A$ and $A \subseteq A \cap B$, we have that $A \cap B = A$. □

[1] See Definition 3.1.7.

[2] See Definitions 3.1.3 and 3.1.9.

Now we present a few more examples of set theorems and proofs, in order to set you up for the exercises to follow.

Theorem 3.2.4 (DeMorgan's Law). If A and B are sets, then

$$(A \cap B)^C = A^C \cup B^C \tag{3.8}$$

Proof.

($\subseteq$): Pick $x \in (A \cap B)^C$. Then $x \notin A \cap B$, so that either $x \notin A$ or $x \notin B$. We consider each case.

(Case $x \notin A$): If $x \notin A$, then $x \in A^C$. Thus $x \in A^C \cup B^C$.

(Case $x \notin B$): If $x \notin B$, then $x \in B^C$. Thus $x \in A^C \cup B^C$.

In either case, we have that $x \in A^C \cup B^C$, so that $(A \cap B)^C \subseteq A^C \cup B^C$.

($\supseteq$): Pick $x \in A^C \cup B^C$. Then either $x \in A^C$ or $x \in B^C$. We consider each case.

(Case $x \in A^C$): If $x \in A^C$, then $x \notin A$. But if $x \notin A$, then certainly $x \notin A \cap B$. Thus $x \in (A \cap B)^C$.

(Case $x \in B^C$): If $x \in B^C$, then $x \notin B$. But if $x \notin B$, then certainly $x \notin A \cap B$. Thus $x \in (A \cap B)^C$.

In either case, we have that $x \in (A \cap B)^C$, so that $A^C \cup B^C \subseteq (A \cap B)^C$.

Since $(A \cap B)^C \subseteq A^C \cup B^C$ and $(A \cap B)^C \supseteq A^C \cup B^C$, we have that $(A \cap B)^C = A^C \cup B^C$. □

The fact that the empty set contains no elements can make for some interesting twists in proofs.

Theorem 3.2.5 For any set A, $A \cup \emptyset = A$.

Proof.

($\subseteq$): Pick $x \in A \cup \emptyset$. Then either $x \in A$, or $x \in \emptyset$. But since $\emptyset$ contains no elements, it must be that $x \in A$. Thus $A \cup \emptyset \subseteq A$.

($\supseteq$): Clearly, if $x \in A$, then $x \in A \cup \emptyset$. Thus $A \subseteq A \cup \emptyset$.

Therefore, $A \cup \emptyset = A$. □

The next theorem is natural to prove by contrapositive.

Theorem 3.2.6 If $A \subseteq B$, then $A - B = \emptyset$.

Proof. Suppose $A - B$ is non-empty. Then there exists $x \in (A - B)$. Thus $x \in A \cap B^C$, so that $x \in A$ and $x \in B^C$. But the existence of such an x is precisely the negation of Definition 3.1.7, so that $A \not\subseteq B$. □

Proof by contradiction can come in handy, too.

Theorem 3.2.7 If A is any set, $\emptyset \subseteq A$.

Proof. Suppose there exists a set A such that $\emptyset \not\subseteq A$. Then there exists $x \in \emptyset$ such that $x \notin A$. But $\emptyset$ contains no elements. This contradicts the definition of $\emptyset$. Thus $\emptyset \subseteq A$. □

Exercise 3.2.8 State the converse of Theorem 3.2.1 and prove it.

Exercise 3.2.9 Prove the following.

(a) For all sets A and B, if $A \subseteq B$, then $A \cup B = B$.

(b) If $A \subseteq B$ and $B \subseteq C$, then $A \subseteq C$.

(c) $A \cap A^C = \emptyset$.

(d) If $A \subseteq B$, then $A^C \supseteq B^C$.

(e) (DeMorgan's Law) For all sets A and B, $(A \cup B)^C = A^C \cap B^C$.

(f) If A and B are disjoint, then $A \Delta B = A \cup B$.

(g) If $A \subseteq C$ and $B \subseteq C$, then $A \cup B \subseteq C$.

(h) If $C \subseteq A$ and $C \subseteq B$, then $C \subseteq A \cap B$.

(i) If $A \subseteq B$ and $C \subseteq D$, then $A \cap C \subseteq B \cap D$.

(j) If $A \subseteq B$ and $C \subseteq D$, then $A \cup C \subseteq B \cup D$.

Exercise 3.2.10 State and prove the converse of Exercise 3.2.9(f).[3]

Exercise 3.2.11 Prove that $\cap$ distributes over $\cup$ and vice versa.

Exercise 3.2.12 Sometimes we can prove certain sets are equal without having to chase elements back and forth, by appealing to earlier theorems we have proved. By making appropriate references to certain results from Exercise 3.2.9, prove the following are true for all sets M, N, S, and T.

(a) If $M = N$ and $S = T$, then $M \cap S = N \cap T$.

(b) If $M = N$ and $S = T$, then $M \cup S = N \cup T$.

(c) If $M = N$, then $M^C = N^C$.

[3] Try proof by contrapositive or contradiction.

Notice that Exercise 3.2.12 says that set intersection, union, and complement are well-defined operations on sets.

Exercise 3.2.13 Show that union and intersection do not allow cancellation by providing counterexamples to the following statements.

(a) If $A \cup C = B \cup C$, then $A = B$.

(b) If $A \cap C = B \cap C$, then $A = B$.

Exercise 3.2.14 If X and Y are disjoint sets, we sometimes write $X \cup Y$ as $X \dot{\cup} Y$. This is a way of talking about the set $X \cup Y$ by tagging it with a little symbol (the dot) that tells the reader the additional information that X and Y are disjoint. So if you see a statement like

$$A \cup B = A \dot{\cup} (B - A) \tag{3.9}$$

it is equivalent to the compound statement

$$A \cup B = A \cup (B - A) \quad \text{and} \quad A \cap (B - A) = \emptyset \tag{3.10}$$

Prove Eq. (3.9) by showing both parts of (3.10).

3.3 Families of Sets

If you're working with only a few sets at a time as we did in Section 3.2, it's probably sufficient to use A, B, C, and so on, to represent them. If you have a set of, say, 10 sets (generally called a *family* or *collection* of sets instead of a set of sets), it might be more sensible to put them into a family and address them as $A_1, A_2, \ldots, A_{10}$. In a case like this, we would say that the set $\{1, 2, 3, \ldots, 10\}$ *indexes* the family of sets. If we write

$$\mathbb{N}_n = \{1, 2, 3, \ldots, n\}$$

then we could write a family of n sets as

$$\{A_1, A_2, A_3, \ldots, A_n\} = \{A_k : k \in \mathbb{N}_n\} = \{A_k\}_{k=1}^n$$

and we would say that $\mathbb{N}_n$ is an *index set* for the family of sets. This notation has advantages, for then we could write unions and intersections more succinctly:

$$A_1 \cup A_2 \cup A_3 \cup \cdots \cup A_n = \bigcup_{k=1}^{n} A_k \tag{3.11}$$

$$A_1 \cap A_2 \cap A_3 \cap \cdots \cap A_n = \bigcap_{k=1}^{n} A_k \tag{3.12}$$

In Definitions 3.3.3 and 3.3.4, we will define precisely what we mean by this sort of union and intersection.

We can go even further. It is conceivable we might need to work with infinitely many sets $\{A_1, A_2, A_3, \dots\}$ that we might want to index with the positive integers. For example, if we use the familiar interval notation

$$[a, b] = \{x : x \in \mathbb{R} \text{ and } a \leq x \leq b\}$$

we might talk about the family of intervals $\mathcal{F} = \{A_n\}_{n\in\mathbb{N}}$, where $A_n = [0, 1/n]$. To form the union or intersection of a family of sets indexed by the positive integers, we could use notation like that in Eqs. (3.11) and (3.12).

$$\bigcup_{n=1}^{\infty} A_n \quad \text{and} \quad \bigcap_{n=1}^{\infty} A_n \tag{3.13}$$

or we could write something like

$$\bigcup_{n\in\mathbb{N}} A_n \quad \text{and} \quad \bigcap_{n\in\mathbb{N}} A_n \tag{3.14}$$

where by (3.14) we understand that n is allowed to take on all values of the indexing set of natural numbers.

The notation in (3.14) is handy when the indexing set is more complicated than the positive integers and does not allow us to think of some index variable n starting at 1 and progressing sequentially off to infinity. For it is conceivable that any set $\mathcal{A}$ can index a family of sets. We can then address individual sets in the family as A_α, where $\alpha \in \mathcal{A}$, and denote the family $\mathcal{F} = \{A_\alpha\}_{\alpha\in\mathcal{A}}$. Union and intersection could then be written as

$$\bigcup_{\alpha\in\mathcal{A}} A_\alpha \quad \text{and} \quad \bigcap_{\alpha\in\mathcal{A}} A_\alpha$$

Example 3.3.1 One important contribution of the German mathematician Richard Dedekind (1831–1916) is a rigorous foundation of the set of real numbers. He employed what is now called a *Dedekind cut*, whereby the set of rational numbers is “cut” into two pieces. Specifically, for a real number r, Dedekind worked with sets of the form

$$A_r = \{x : x \in \mathbb{Q} \text{ and } x < r\} \tag{3.15}$$

$$B_r = \{x : x \in \mathbb{Q} \text{ and } x > r\} \tag{3.16}$$

For example, although $\sqrt{2}$ is irrational (as you will see in Section 3.10), it certainly makes sense to talk about $A_{\sqrt{2}}$, the set of all rational numbers less than $\sqrt{2}$. The real numbers form the index set for the families $\{A_r\}_{r\in\mathbb{R}}$ and $\{B_r\}_{r\in\mathbb{R}}$. ■

Example 3.3.2 Consider the family of intervals

$$\mathcal{F} = \{[-r, r] : r \in \mathbb{R}^+\} \tag{3.17}$$

where $\mathbb{R}^+$ denotes the positive real numbers. The index set is $\mathbb{R}^+$, and we might want to use the notation $I_r = [-r, r]$ to represent one interval in the family. We might also write $\mathcal{F} = \{I_r\}_{r\in\mathbb{R}^+}$. ■

Of course, we don't have to worry about having our family of sets so well organized as to be indexed by any set at all. In the same way that we talk about a set A and address an arbitrary element $x \in A$, we can call our family of sets $\mathcal{F}$ and address an arbitrary set in the family as $A \in \mathcal{F}$. Then the union and intersection of the sets in $\mathcal{F}$ could be written as

$$\bigcup_{A\in\mathcal{F}} A \quad \text{and} \quad \bigcap_{A\in\mathcal{F}} A \tag{3.18}$$

or more simply as

$$\bigcup_{\mathcal{F}} A \quad \text{and} \quad \bigcap_{\mathcal{F}} A \tag{3.19}$$

The statements in (3.18) and (3.19) assume the least structure on the family of sets, so they are a good general notation.

Before we dig into families of sets, we're going to construct an extended metaphor that will help carry us through the complexities of the ideas and notation. Families of sets tend to confuse students at first, and for several reasons. First, when we talk about a family, we are actually talking about a set whose elements are sets. So it might be that $x \in A$ and $A \in \mathcal{F}$, but writing $x \in \mathcal{F}$ is wrong, for x is not an element of $\mathcal{F}$, but an element of an element of $\mathcal{F}$. Second, when a family is indexed, the indexing set $\mathcal{A}$ is yet another set to contend with, whose elements α serve as something like name tags by which sets in the family are addressed. So let's stage a little production whose cast of characters all represent sets and terms we have mentioned in this section. Then as we state theorems, the metaphor will help us understand what is being said and how to approach the proof.

Suppose we are studying the mathematics section of the catalog of Prestigious University. Every mathematics class has a number, so let

$$\mathcal{A} = \{104, 210, 211, 212, 330, 360, 430, 431, 510, 531\} \tag{3.20}$$

and think of $\mathcal{A}$ as the set of all the course numbers of mathematics courses offered at PU. This set of course numbers will be our index set. We address an arbitrarily chosen course number as $\alpha \in \mathcal{A}$. Let A_α be the roster of all PU graduates who passed course number α while they were students at PU. We will address an arbitrarily chosen student on roster A_α as x. With these role assignments,

$\mathcal{F} = \{A_\alpha\}_{\alpha\in\mathcal{A}}$ is the set of all these rosters of students who passed the individual courses. That is,

$$\mathcal{F} = \{A_{104}, A_{210}, A_{211}, \ldots, A_{531}\} \tag{3.21}$$

Writing the family of rosters as $\mathcal{F} = \{A_\alpha\}_{\alpha\in\mathcal{A}}$ references each roster in the family in terms of the course number that tags it. Writing $\mathcal{F}$ simply as $\{A\}$ dispenses with the number tags and addresses a particular roster in $\mathcal{F}$ simply as A.

Certainly, the family of sets $\mathcal{F}$ and its indexing set $\mathcal{A}$ need not look anything like the ones we have created here. But thinking of $\mathcal{A}$ as in Eq. (3.20) and $\mathcal{F}$ as in Eq. (3.21) should not limit us or make our proofs less than generally applicable if we use the metaphor as a tool to help clarify our thinking. Just remember that we will address sets in the family in one of two ways. To illustrate the first way, sometimes we will pick some arbitrary $\alpha \in \mathcal{A}$ in order to talk about A_α, which is like choosing an arbitrary course number and talking about the roster for the course with that number. Or perhaps we will claim the existence of some specific $\alpha_0 \in \mathcal{A}$ in order to talk about A_{α_0}. In our metaphor, this is the claim that there is a certain course number whose roster has some property. To illustrate the second way of addressing sets in the family, sometimes we might pick an arbitrary $A \in \mathcal{F}$, which is like choosing an arbitrary mathematics roster without any specific reference to a course number. Or perhaps we will claim the existence of some specific $A_0 \in \mathcal{F}$. This is the claim that there is a particular mathematics course roster that has a certain property, without making any reference to the course number.

Having created this little metaphor as an aid to our understanding, let's define more terms and derive some results, using the metaphor to motivate and get us over some humps. First, how should we define the sets $\cup_{A\in\mathcal{F}}A$ and $\cap_{A\in\mathcal{F}}A$? Or equivalently, how should we define the following statements?[4]

$$x \in \bigcup_{\mathcal{F}} A \quad \text{and} \quad x \in \bigcap_{\mathcal{F}} A \tag{3.22}$$

We want $\cup_{A\in\mathcal{F}}A$ to be the set you get when you take each A in $\mathcal{F}$ and dump all its elements into a single set. For our metaphor, $\cup_{\alpha\in\mathcal{A}}A_\alpha$, or $\cup_{A\in\mathcal{F}}A$, however you choose to write it, is the single roster of graduates created by unioning all the individual mathematics class rosters. Maybe you can see that $\cup_{\mathcal{F}}A$ consists of all graduates who have ever passed a mathematics class at PU. So if x represents a graduate, saying $x \in \cup_{\mathcal{F}}A$ means that there is some mathematics course at PU that x passed. We can write this in two ways, with or without a reference to the course numbers:

$$x \in \bigcup_{\alpha\in\mathcal{A}} A_\alpha \Leftrightarrow (\exists\alpha \in \mathcal{A})(x \in A_\alpha) \tag{3.23}$$

[4] Think in terms of ∃ for the union and ∀ for the intersection.

$$x \in \bigcup_{A \in \mathcal{F}} A \Leftrightarrow (\exists A \in \mathcal{F})(x \in A) \tag{3.24}$$

The form of (3.23) says there exists a course number $\alpha \in \mathcal{A}$ such that graduate x is on the roster of mathematics course number α. The form of (3.24) says simply that there is a mathematics course roster on which the name of graduate x appears, without any reference to a course number. With this, we arrive at the following definition.

Definition 3.3.3 Let $\mathcal{F}$ be a family of sets indexed by $\mathcal{A}$. Then the *union over* $\mathcal{F}$ is defined by

$$\bigcup_{\alpha \in \mathcal{A}} A_\alpha = \{x : (\exists \alpha \in \mathcal{A})(x \in A_\alpha)\} \tag{3.25}$$

Without reference to the index set, this becomes

$$\bigcup_{A \in \mathcal{F}} A = \{x : (\exists A \in \mathcal{F})(x \in A)\} \tag{3.26}$$

Now what does $\cap_{\alpha \in \mathcal{A}} A_\alpha$ (or $\cap_{A \in \mathcal{F}} A$) correspond to in our metaphor? This is the intersection of all rosters, so it is the list of all graduates who passed all mathematics courses at PU. So what does it mean mathematically to say $x \in \cap_{\alpha \in \mathcal{A}} A_\alpha$ (or $x \in \cap_{\mathcal{F}} A$)?

$$x \in \bigcap_{\alpha \in \mathcal{A}} A_\alpha \Leftrightarrow (\forall \alpha \in \mathcal{A})(x \in A_\alpha) \tag{3.27}$$

$$x \in \bigcap_{A \in \mathcal{F}} A \Leftrightarrow (\forall A \in \mathcal{F})(x \in A) \tag{3.28}$$

With this we arrive at the following definition.

Definition 3.3.4 Let $\mathcal{F}$ be a family of sets indexed by $\mathcal{A}$. Then the *intersection over* $\mathcal{F}$ is defined by

$$\bigcap_{\alpha \in \mathcal{A}} A_\alpha = \{x : (\forall \alpha \in \mathcal{A})(x \in A_\alpha)\} \tag{3.29}$$

Without reference to the index set, this becomes

$$\bigcap_{A \in \mathcal{F}} A = \{x : (\forall A \in \mathcal{F})(x \in A)\} \tag{3.30}$$

Many of the results we proved in Section 3.2 for a family of two or three sets carry over to analogous results for a family of sets of any size. So as you work your way through the rest of this section, note that some of the theorems are generalizations of results from Section 3.2.

Example 3.3.5 Construct the negation of the statement $x \in \cup_{\mathcal{F}} A$.

Solution By negating the statement in (3.23), to say $x \notin \cup_{\mathcal{F}} A$ means

$$(\forall \alpha \in \mathcal{A})(x \notin A_\alpha)$$

Without making reference to the index set, we use (3.24) to have

$$(\forall A \in \mathcal{F})(x \notin A) \qquad \blacksquare$$

In our mathematics class metaphor, what does $x \notin \cup_{\mathcal{F}} A$ mean? Graduate x is not on the universal roster of mathematics classes, so x never passed a mathematics class at PU. That is, for every course number $\alpha \in \mathcal{A}$, x did not pass the mathematics course numbered α, or, for every roster $A \in \mathcal{F}$, x is not on roster A.

EXERCISE 3.3.6 Construct the negation of the statement $x \in \cap_{\mathcal{F}} A$ in two forms. What does $x \notin \cap_{\mathcal{F}} A$ mean in the context of the mathematics class metaphor?

Theorem 3.3.7 (DeMorgan's Law). Suppose $\mathcal{F}$ is a family of sets. Then

$$\left[\bigcup_{\mathcal{F}} A\right]^C = \bigcap_{\mathcal{F}} A^C \tag{3.31}$$

Proof.

($\subseteq$): Pick $x \in [\cup_{\mathcal{F}} A]^C$. Then $x \notin \cup_{\mathcal{F}} A$. Therefore, for all $A \in \mathcal{F}$, $x \notin A$. But then $x \in A^C$ for every $A \in \mathcal{F}$. Thus $x \in \cap_{\mathcal{F}} A^C$.

($\supseteq$): Pick $x \in \cap_{\mathcal{F}} A^C$. Then $x \in A^C$ for every $A \in \mathcal{F}$. Thus, $x \notin A$ for every $A \in \mathcal{F}$, so that $x \notin \cup_{\mathcal{F}} A$. Therefore $x \in [\cup_{\mathcal{F}} A]^C$. □

To see what Theorem 3.3.7 says in our metaphor, we need a universal set in which the sets in the family are meaningful. Let U be the roster of all PU graduates. With this, what set is being talked about in Theorem 3.3.7, and what are the two ways of constructing it in Eq. (3.31)? To construct the left-hand side of Eq. (3.31), we first combine all the mathematics rosters into a single roster and then take all the PU graduates *except* these. This is the list of all PU graduates who avoided mathematics altogether while they were at PU. To construct the right-hand side of Eq. (3.31), first we take each mathematics class roster and consider the complement. For example, A_{211}^C is the list of all PU graduates who did not pass Mathematics 211. By taking the intersection of all these complements, we arrive at the list of PU graduates who avoided Math 104 *and* Math 210 *and* ... *and* Math 531; that is, the list of graduates who avoided mathematics altogether.

EXERCISE 3.3.8 [DeMorgan's Law] Suppose $\mathcal{F}$ is a family of sets. Then

$$\left[\bigcap_{\mathcal{F}} A\right]^C = \bigcup_{\mathcal{F}} A^C \tag{3.32}$$

For the next theorem, we prove part 2 here and leave part 1 to you as an exercise.

Theorem 3.3.9 Let $\mathcal{F}$ be a family of sets indexed by $\mathcal{A}$, and suppose $\mathcal{B} \subseteq \mathcal{A}$. Then,

1. $\cup_{\beta\in\mathcal{B}} A_\beta \subseteq \cup_{\alpha\in\mathcal{A}} A_\alpha$
2. $\cap_{\beta\in\mathcal{B}} A_\beta \supseteq \cap_{\alpha\in\mathcal{A}} A_\alpha$

Before we prove part 2 of Theorem 3.3.9, let's see how it relates to our metaphor. Since $\mathcal{A}$ is the set of all mathematics course numbers at PU, it might work to think of $\mathcal{B}$ as the set of all course numbers of lower level mathematics courses at PU. So

$$\mathcal{B} = \{104, 210, 211, 212\} \tag{3.33}$$

With that, what does part 2 of Theorem 3.3.9 say? The set construction on the left-hand side is the intersection of the class rosters across all the lower level courses, while the right-hand side is the intersection of the rosters of all the mathematics classes. So if we pick some $x \in \cap_{\alpha\in\mathcal{A}} A_\alpha$, then x is a PU graduate who passed every mathematics course offered. So certainly x passed all the lower level courses. Here is the proof.

Proof of part 2. Pick $x \in \cap_{\alpha\in\mathcal{A}} A_\alpha$. We must show that $x \in A_\beta$ for all $\beta \in \mathcal{B}$, so pick $\beta \in \mathcal{B}$. Since $\mathcal{B} \subseteq \mathcal{A}$, it follows that $\beta \in \mathcal{A}$. Therefore $x \in A_\beta$, because $x \in A_\alpha$ for any $\alpha \in \mathcal{A}$. Since β was chosen arbitrarily, we have shown that $x \in A_\beta$ for all $\beta \in \mathcal{B}$, so that $x \in \cap_{\beta\in\mathcal{B}} A_\beta$. □

EXERCISE 3.3.10 Prove part 1 of Theorem 3.3.9. What is this statement saying in the context of the mathematics class metaphor?

We can write Theorem 3.3.9 in a slightly different form if the family of sets is not indexed. If $\mathcal{F}_1$ is a family of sets and $\mathcal{F}_2 \subseteq \mathcal{F}_1$, we call $\mathcal{F}_2$ a *subfamily* of $\mathcal{F}_1$. Since $\mathcal{B} \subseteq \mathcal{A}$ in Theorem 3.3.9, $\{A_\beta\}_{\beta\in\mathcal{B}}$ is a subfamily of $\{A_\alpha\}_{\alpha\in\mathcal{A}}$. Swapping the notation in Theorem 3.3.9 for an arbitrary family $\mathcal{F}_1$ and a subfamily $\mathcal{F}_2$, we have the following.

Theorem 3.3.11 Suppose $\mathcal{F}_1$ is a family of sets, and $\mathcal{F}_2$ is a subfamily of $\mathcal{F}_1$.

Then

1. $\cup_{\mathcal{F}_2} A \subseteq \cup_{\mathcal{F}_1} A$
2. $\cap_{\mathcal{F}_2} A \supseteq \cap_{\mathcal{F}_1} A$

The next theorem involves two families of sets, both indexed by $\mathcal{A}$, where corresponding sets in the two families are related by subset inclusion. To understand what it says, think of B_α as the set of all female PU graduates who passed mathematics course number α.

Exercise 3.3.12 Suppose $\mathcal{F}_1 = \{A_\alpha\}_{\alpha \in \mathcal{A}}$ and $\mathcal{F}_2 = \{B_\alpha\}_{\alpha \in \mathcal{A}}$ are two families of sets with the property that $B_\alpha \subseteq A_\alpha$ for every $\alpha \in \mathcal{A}$. Then

(a) $\bigcup_{\alpha \in \mathcal{A}} B_\alpha \subseteq \bigcup_{\alpha \in \mathcal{A}} A_\alpha$

(b) $\bigcap_{\alpha \in \mathcal{A}} B_\alpha \subseteq \bigcap_{\alpha \in \mathcal{A}} A_\alpha$

Exercise 3.3.13 Suppose $\mathcal{F} = \{A\}$ is a family of sets, and suppose C is a set for which $A \subseteq C$ for every $A \in \mathcal{F}$. Show that $\cup_{\mathcal{F}} A \subseteq C$.[5]

Exercise 3.3.14 Suppose $\mathcal{F} = \{A\}$ is a family of sets, and suppose D is a set for which $D \subseteq A$ for every $A \in \mathcal{F}$. Show that $D \subseteq \cap_{\mathcal{F}} A$.[6]

3.4 The Principle of Mathematical Induction

In the world of mathematics, the well-ordering principle (WOP) is often taken as an axiom. In this section, we derive a theorem based on the WOP called the *Principle of Mathematical Induction* (PMI). To write a *proof by induction*, the imagery is that we have an infinite row of dominoes that we must knock down. First, we show figuratively that we can knock the first domino down. Then we show that if the nth domino falls, then so does the $(n+1)$st. This very important proof technique is useful when the theorem you're trying to prove has a form like any of the following examples.

Example 3.4.1 For any positive integer n,

$$1 + 2 + 3 + \cdots + n = \sum_{k=1}^{n} k = \frac{n(n+1)}{2} \tag{3.34}$$

■

[5] Think of C as the set of all PU graduates who ever enrolled in a mathematics class.

[6] Think of D as the set of all PU mathematics majors (whom we will assume would have taken every mathematics course) who graduated with a 4.00 GPA.

Example 3.4.2 For any positive integer n,

$$1^2 + 2^2 + 3^2 + \cdots + n^2 = \sum_{k=1}^{n} k^2 = \frac{n(n+1)(2n+1)}{6} \tag{3.35}$$

■

Example 3.4.3 Let A be a given set, and $\{B_k\}_{k=1}^{n}$ a family of $n \geq 1$ sets. Then

$$A \cup (B_1 \cap B_2 \cap \cdots \cap B_n) = (A \cup B_1) \cap (A \cup B_2) \cap \cdots \cap (A \cup B_n) \tag{3.36}$$

■

Example 3.4.4 Suppose there are $n \geq 1$ people in a room, and everyone shakes hands with everyone else exactly once. Then the number of handshakes that takes place is $n(n-1)/2$. ■

Example 3.4.5 Let G be a connected, acyclic graph on n vertices, where $n \geq 1$. Then G has $n - 1$ edges. ■

Notice that Examples 3.4.1–3.4.5 all make a statement about a finite but unspecified n number of things, and you want to prove that the claim is true for any $n \geq 1$.

You might have found Eq. (3.34) handy if you had been in grammar school with Carl Friedrich Gauss in the 1780s. Gauss, a very precocious child, showed amazing mathematical ability at a very early age. A somewhat embellished story goes that when Gauss was eight years old, it was raining one day during recess, Internet access was down, and his teacher needed to keep the children in the class busy for a while. So he told them to add up the first 100 natural numbers without their calculators. Gauss figured out how to get the result quickly in the following way. By writing the sum twice, once in reverse order, he added vertically, term by term:

$$\begin{array}{r} 1+\quad 2+\quad 3+\quad 4+\cdots+\quad 99+100 \\ 100+\ 99+\ 98+\ 97+\cdots+\quad 2+\quad 1 \\ \hline \underbrace{101+101+101+101+\cdots+101+101}_{100\text{ terms}} \end{array}$$

Gauss observed that 100 × 101 is twice the desired result, so he quickly reported the result of 5050. If you perform a similar trick replacing 100 with an arbitrary n, you get Eq. (3.34).

Though Gauss's technique might seem sufficient as a proof, there is something a little sloppy about making a claim that involves a "dot dot dot" in it. The PMI is a theorem derived from the WOP that eliminates this untidiness.

So what is the PMI? And what does the WOP have to do with it? Let's conduct a thought experiment to set it up, first in its standard form. Then in Section 3.5 we will look at some variations.

Suppose we consider a set S, which is assumed to be a subset of the positive integers, and suppose S is known to have the following properties:

(I 1) $1 \in S$

(I 2) If $n \geq 1$ and $n \in S$, then $n + 1 \in S$.

Then what precisely is S? Now I 1 says $1 \in S$, but then I 2 applies to guarantee that, since $1 \in S$, then $1 + 1 = 2 \in S$. But then I 2 applies again to guarantee that $2 + 1 = 3 \in S$, and so on. Now $S \subseteq \mathbb{N}$ is assumed, and it appears that every positive integer is also in S, so that $\mathbb{N} \subseteq S$. So surely $S = \mathbb{N}$. Fair enough. But this argument has its own "dot dot dot" and is a bit less than rigorous.

Another way to look at this same argument is still a little open ended, but comes closer to the actual proof as we will present it. If it were true that $S \neq \mathbb{N}$, then since $S \subseteq \mathbb{N}$, it must be that $\mathbb{N} \not\subseteq S$. Thus there exists $k \in \mathbb{N}$ such that $k \notin S$. By I 1, $k \neq 1$, so that $k - 1$ is a positive integer. By the contrapositive of I 2, since $k \notin S$, then $k - 1 \notin S$ either. Therefore $k - 1 \neq 1$, so that $k - 2 \in \mathbb{N}$, and since $k - 1 \notin S$, I 2 implies $k - 2 \notin S$. Here's where the "dot dot dot" comes in, and we have that $k, k - 1, k - 2, \ldots \notin S$. But this crashes head-on into the fact that $1 \in S$.

If we look to the WOP, then we can clean up these arguments and find in the WOP enough strength to provide a rigorous proof that any subset of the positive integers that has properties I 1–I 2 must, in fact, be the entire set of positive integers. This is the Principle of Mathematical Induction.

Theorem 3.4.6 (PMI). Suppose S is a subset of the positive integers with the following properties:

(I 1) $1 \in S$

(I 2) If $n \geq 1$ and $n \in S$, then $n + 1 \in S$.

Then $S = \mathbb{N}$.

The proof of the PMI is a great example of proof by contradiction. If you would like to try to prove it yourself, then take a glance at the hints[7,8,9,10,11] below if you need to. Here is the proof.

Proof. Suppose S is a subset of the positive integers with properties I 1–I 2. Suppose also that $S \neq \mathbb{N}$. Then $\mathbb{N} \not\subseteq S$. Thus there exists $n \in \mathbb{N}$ such that $n \notin S$. If we

[7] The theorem says $[(S \subseteq \mathbb{N}) \wedge \text{I}\,1 \wedge \text{I}\,2] \to (S = \mathbb{N})$. Suppose this is false.

[8] State $\mathbb{N} \not\subseteq S$ in $\exists$ form.

[9] Let $T = \mathbb{N} - S$. What does the WOP say about T?

[10] If T has a smallest element a, what can you say about $a - 1$?

[11] But if $a - 1 \in S$, then what is true of a?

define $T = \mathbb{N} - S = \mathbb{N} \cap S^C$, then $T \subseteq \mathbb{N}$ and $n \in T$. Thus T is a non-empty subset of the positive integers, which, by the WOP, contains a smallest element a. Now it is impossible that $a = 1$ because $1 \in S$. Thus $a > 1$, so that $a - 1 \in \mathbb{N}$. Furthermore, since a is the smallest element of T, it must be that $a - 1 \notin T$. Thus $a - 1 \in S$. But by I2, since $a - 1 \in S$, it follows that $a - 1 + 1 = a \in S$. But if $a \in S$, then $a \notin T$. This is a contradiction. Thus there is no smallest element of T, which means that T is empty. Therefore, $\mathbb{N} \subseteq S$, whence $S = \mathbb{N}$. □

What does Theorem 3.4.6 have to do with proving the claims in Examples 3.4.1–3.4.5? Think of these examples as statements about the positive integers. They say effectively that a certain formula or statement is true for all $n \geq 1$. To be as general as possible, address such a statement by $P(n)$. Let's define S to be the set of all positive integers n for which the statement $P(n)$ is true. The trick is to show that S has properties I1–I2. Then the PMI will allow us to conclude that $S = \mathbb{N}$, which is the same as saying $P(n)$ is true for all n. First we show that $1 \in S$ by showing $P(1)$ is true. This is called establishing the *base case*. Then we show that $n \in S \Rightarrow n + 1 \in S$ by supposing $P(n)$ is true, and using this to show that $P(n + 1)$ is true. The assumption that $n \in S$ is called the *inductive assumption*, and the part of the proof where we show that $n \in S$ implies $n + 1 \in S$ is called the *inductive step*. Having shown that S has properties I1–I2, we then conclude $S = \mathbb{N}$; that is, $P(n)$ holds true for all n. Here is a sample theorem.

Theorem 3.4.7 (Sample). For all positive integers n, $P(n)$.

Proof. We prove by induction on $n \geq 1$.

(I1) $P(1)$.

(I2) Suppose $n \geq 1$ and $P(n)$. Then .... Thus $P(n + 1)$.

Therefore, by induction, $P(n)$ is true for all $n \geq 1$. □

There will be some point in step I2 where you use the inductive assumption $P(n)$ to get over the hump of showing $P(n + 1)$. In the proof of the next theorem, we rewrite a sum in the following way:

$$\sum_{k=1}^{n+1} a_k = (a_1 + a_2 + \cdots + a_n) + a_{n+1} = \left(\sum_{k=1}^{n} a_k\right) + a_{n+1} \tag{3.37}$$

Theorem 3.4.8 For all $n \geq 1$,

$$\sum_{k=1}^{n} k^2 = \frac{n(n + 1)(2n + 1)}{6} \tag{3.38}$$

Proof. We prove by induction on $n \geq 1$.

(I1) For the case $n = 1$, we have

$$\sum_{k=1}^{n} k^2 = \sum_{k=1}^{1} k^2 = 1^2 = 1 = 1(1+1)(2 \cdot 1 + 1)/6 = n(n+1)(2n+1)/6$$

so that Eq. (3.38) holds true for $n = 1$.

(I2) Suppose $n \geq 1$ and $\sum_{k=1}^{n} k^2 = n(n+1)(2n+1)/6$. Then

$$\begin{aligned}\sum_{k=1}^{n+1} k^2 &= \sum_{k=1}^{n} k^2 + (n+1)^2 = \frac{n(n+1)(2n+1)}{6} + (n+1)^2 \\ &= \frac{2n^3 + 3n^2 + n + 6(n^2 + 2n + 1)}{6} = \frac{2n^3 + 9n^2 + 13n + 6}{6} \\ &= \frac{(n+1)(n+2)(2n+3)}{6} = \frac{(n+1)[(n+1)+1][2(n+1)+1]}{6}\end{aligned}$$

Thus Eq. (3.38) holds for $n + 1$.

By the PMI, it follows that Eq. (3.38) holds for all $n \geq 1$. □

EXERCISE 3.4.9 Prove the following sum formulas for $n \geq 1$.

(a) $\sum_{k=1}^{n} k = \frac{n(n+1)}{2}$

(b) $\sum_{k=1}^{n} (-1)^k k^2 = (-1)^n \frac{n(n+1)}{2}$

(c) $\sum_{k=1}^{n} (2k - 1) = n^2$

(d) $\sum_{k=1}^{n} k^3 = \left(\sum_{k=1}^{n} k\right)^2$

(e) $\sum_{k=1}^{n} \frac{1}{k(k+1)} = \frac{n}{n+1}$

Before we prove the claim in Example 3.4.3, we need to make an observation. In Exercise 3.2.11, you showed that

$$A \cup (B \cap C) = (A \cup B) \cap (A \cup C) \tag{3.39}$$

This is the special case of Eq. (3.36) where $n = 2$. We need this fact in proving the inductive step below.

Theorem 3.4.10 Suppose A is a given set and $\{B_k\}_{k=1}^{n}$ is a family of $n \geq 1$ sets. Then

$$A \cup (B_1 \cap B_2 \cap \cdots \cap B_n) = (A \cup B_1) \cap (A \cup B_2) \cap \cdots \cap (A \cup B_n) \tag{3.40}$$

Proof. To clean up the notation, we write Eq. (3.40) as

$$A \cup [\cap_{k=1}^{n} B_k] = \cap_{k=1}^{n} (A \cup B_k) \tag{3.41}$$

(I1) If $n = 1$, then $\cap_{k=1}^{1} B_k = B_1$, so that both sides of Eq. (3.41) are simply $A \cup B_1$.

(I2) Suppose $n \geq 1$ and Eq. (3.41) holds for n. Then

$$\begin{aligned} A \cup [\cap_{k=1}^{n+1} B_k] &= A \cup [(\cap_{k=1}^{n} B_k) \ \cap \ B_{n+1}] \\ &\overset{\text{by (3.39)}}{=} [A \cup (\cap_{k=1}^{n} B_k)] \cap (A \cup B_{n+1}) \\ &= [\ \cap_{k=1}^{n} (A \cup B_k)] \cap (A \cup B_{n+1}) = \cap_{k=1}^{n+1} (A \cup B_k) \end{aligned}$$

Thus, Eq. (3.41) holds for $n + 1$.

By the PMI, Eq. (3.41) holds for all n. □

EXERCISE 3.4.11 Suppose A is a given set and $\{B_k\}_{k=1}^{n}$ is a family of $n \geq 1$ sets. Then

$$A \cap (B_1 \cup B_2 \cup \cdots \cup B_n) = (A \cap B_1) \cup (A \cap B_2) \cup \cdots \cup (A \cap B_n) \tag{3.42}$$

EXERCISE 3.4.12 Suppose $A_1, A_2, A_3, \ldots$ are sets with the property that $A_n \subseteq A_{n+1}$ for every $n \geq 1$. Use induction to show that $A_1 \subseteq A_n$ for all n.

Now we prove the claim in Example 3.4.4.

Theorem 3.4.13 Suppose there are n people in a room, and everyone shakes hands with everyone else exactly once. Then the number of handshakes that takes place is $n(n-1)/2$.

Proof.

(I1) If there is one person in the room, then there are zero handshakes. But $1(1-1)/2 = 0$, so the result is true for the case $n = 1$.

(I2) Suppose $n \geq 1$ and that in any room of n people there will be $n(n-1)/2$ handshakes. Now suppose there are $n + 1$ people in the room. Remove one person from the room, so that n people remain. If all these people shake hands with each other, then by the inductive assumption, there will be $n(n-1)/2$ handshakes. Now bring the removed person back into the room. When this person shakes hands with everyone, there will be n more handshakes. Thus the total number of handshakes among all $n + 1$ people is

$$\frac{n(n-1)}{2} + n = \frac{n^2 - n}{2} + \frac{2n}{2} = \frac{n^2 + n}{2} = \frac{(n+1)n}{2}$$

Thus the result is true for the group of $n+1$ people.

By the PMI, it follows that the result is true for all $n \geq 1$. □

Exercise 3.4.14 The Tower of Hanoi game consists of three pegs, one of which holds a stack of n disks, and where the disks decrease in size from the bottom to the top of the stack. See Figure 3.9. The object of the game is to move the entire stack of disks from its current peg to either of the other two pegs. The rules of the game are that the disks may be moved from peg to peg only one at a time, and at no point may a disk sit on top of a smaller disk. Show that the object of the game may be achieved on a stack of $n \geq 1$ disks in $2^n - 1$ moves.

Exercise 3.4.15 Suppose n is a positive integer and $a, b_1, b_2, \ldots, b_n$ are real numbers. Then

$$a\left(\sum_{k=1}^{n} b_k\right) = \sum_{k=1}^{n}(ab_k) \tag{3.43}$$

Theorem 3.4.16 Suppose m and n are positive integers and let $a_1, a_2, \ldots, a_m$ and $b_1, b_2, \ldots, b_n$ be real numbers. Then

$$\left(\sum_{j=1}^{m} a_j\right)\left(\sum_{k=1}^{n} b_k\right) = \sum_{j=1}^{m}\left(\sum_{k=1}^{n} a_j b_k\right) \tag{3.44}$$

If $m = n = 2$, this is the FOIL technique of multiplying two binomials.

Proof. First we apply Exercise 3.4.15 by distributing $\sum_{k=1}^{n} b_k$ over the terms in the sum $\sum_{j=1}^{m} a_j$ to have

$$\left(\sum_{j=1}^{m} a_j\right)\left(\sum_{k=1}^{n} b_k\right) = \sum_{j=1}^{m}\left[a_j\left(\sum_{k=1}^{n} b_k\right)\right]$$

Then we apply Exercise 3.4.15 again by distributing each a_j over the terms in the sum $\sum_{k=1}^{n} b_k$ to have

$$\sum_{j=1}^{m}[a_j(\sum_{k=1}^{n} b_k)] = \sum_{j=1}^{m}(\sum_{k=1}^{n} a_j b_k)$$

□

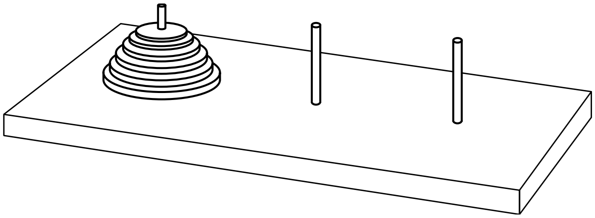

Figure 3.9 The Tower of Hanoi.

Exercise 3.4.17 For real numbers $a_1, a_2, \ldots, a_n$, $\left|\sum_{k=1}^{n} a_k\right| \le \sum_{k=1}^{n} |a_k|$.

Exercise 3.4.18 For all $n \ge 1, 3 \mid (4^n - 1)$.

Exercise 3.4.19 Suppose $a_1, a_2, \ldots, a_n$ are nonzero integers and p is a prime number. Suppose also that $p \mid a_1 a_2 \cdots a_n$. Then there exists some k $(1 \le k \le n)$ such that $p \mid a_k$.

3.5 Variations of the PMI

There are several useful variations of the standard PMI from Section 3.4. First, in defining S on p. 80, there was nothing magical about making 1 the first element of S. We could have supposed that j is any integer and that S is a subset of the integers with the following properties.

(J1) $j \in S$.

(J2) If $n \ge j$ and $n \in S$, then $n + 1 \in S$.

By mimicking almost word for word the proof of Theorem 3.4.6, we have the following.

Theorem 3.5.1 (PMI, Version 2). Suppose S is a subset of the integers that has properties J1–J2. Then $S = \{j, j+1, j+2, \ldots\} = \{n \in \mathbb{Z} : n \ge j\}$.

If $j = 0$, we have a PMI that allows us to prove results are true for all nonnegative integers. The proofs are identical in principle to those in Section 3.4, but the induction arguments are rooted at $n = 0$ instead of $n = 1$.

Exercise 3.5.2 Prove the sum formula $\sum_{k=0}^{n} 2^k = 2^{n+1} - 1$.

Let x be a nonzero real number, and define the following.

$$\begin{aligned} x^0 &= 1 \\ x^{n+1} &= x^n \cdot x \quad \text{for } n \ge 0 \end{aligned} \tag{3.45}$$

This is the standard way of defining exponentiation. Instead of saying something like $x^9 = x \cdots x$ nine times, we say $x^9 = x^8 \cdot x$, and we define nonnegative integer powers of x *recursively*. This definition of exponentiation has some familiar behaviors. Notice that the next theorem deals only with nonnegative exponents. We will prove one result and leave the others to you as an exercise.

Theorem 3.5.3 Let a and b be nonzero real numbers, and let m and n be nonnegative integers. Then

$$a^m \cdot a^n = a^{m+n} \tag{3.46}$$

$$(a^m)^n = a^{mn} \tag{3.47}$$

$$(ab)^n = a^n b^n \tag{3.48}$$

Proof of Eq. (3.46). We prove by induction on $n \geq 0$ (thinking of m as fixed).

(J1) For the case $n = 0$, we have $a^m \cdot a^0 = a^m \cdot 1 = a^m = a^{m+0}$. Thus the result is true for $n = 0$.

(J2) Suppose $n \geq 0$ and $a^m \cdot a^n = a^{m+n}$. We show that $a^m \cdot a^{n+1} = a^{m+(n+1)}$.

$$\begin{aligned} a^m \cdot a^{n+1} &= a^m \cdot (a^n \cdot a) = (a^m \cdot a^n) \cdot a \\ &= a^{m+n} \cdot a = a^{(m+n)+1} = a^{m+(n+1)} \end{aligned}$$

Thus by induction, $a^m \cdot a^n = a^{m+n}$ for all nonzero a and nonnegative integers m and n. □

EXERCISE 3.5.4 Prove Eqs. (3.47) and (3.48).

Now we can define x^n for $n < 0$. For nonzero real number x and positive integer n, define

$$x^{-n} = (x^{-1})^n \tag{3.49}$$

Notice in the next exercise how we must derive the rules for negative exponents very carefully from the same rules for nonnegative exponents. No induction arguments are necessary in this exercise, but merely careful attention to the manipulation of exponents as allowed by Theorem 3.5.3.

EXERCISE 3.5.5 This exercise will show that the rules for exponents in Eqs. (3.46)–(3.48) hold for all integer exponents, positive, zero, or negative. In all parts of this exercise, a and b are nonzero real numbers, and m and n are nonnegative integers. Prove the following.

(a) $a^n \cdot a^{-n} = 1$

(b) $(a^n)^{-1} = (a^{-1})^n$ [12]

(c) $a^{-m} \cdot a^{-n} = a^{-m-n}$

(d) $a^m a^{-n} = a^{m-n}$ [13]

(e) $(a^{-m})^n = a^{-mn}$

(f) $(a^m)^{-n} = a^{-mn}$

[12] This should be a mere observation from part (a).

[13] Take a hint from $a^8 \cdot a^{-6} = a^2 \cdot a^6 \cdot a^{-6} = a^2 = a^{8-6}$ and $a^2 \cdot a^{-8} = a^2 \cdot a^{-2} \cdot a^{-6} = a^{-6} = a^{2-8}$.

(g) $(a^{-m})^{-n} = a^{mn}$

(h) $(ab)^{-n} = a^{-n}b^{-n}$

EXERCISE 3.5.6 Suppose a is a nonzero real number and n is a nonnegative integer. Prove the following.[14]

(a) $(-a)^{2n} = a^{2n}$

(b) $(-a)^{2n+1} = -a^{2n+1}$

EXERCISE 3.5.7 Prove the following for $n \geq 1$.

(a) If $0 \leq a < b$, then $a^n < b^n$ [15]

(b) If $x > 1$, then $x^n > 1$

(c) If $a < b \leq 0$, then

(i) $a^{2n} > b^{2n}$ [16]

(ii) $a^{2n+1} < b^{2n+1}$

(d) If $a < 0 < b$, then $a^{2n+1} < b^{2n+1}$ [17]

Another example of a recursive definition is the *factorial.*

$$\begin{aligned} 0! &= 1 \\ (n+1)! &= (n+1) \cdot n! \quad \text{for } n \geq 0 \end{aligned} \tag{3.50}$$

EXERCISE 3.5.8 Prove the following for $n \geq 0$.

(a) $\sum_{k=0}^{n} \frac{k}{(k+1)!} = 1 - \frac{1}{(n+1)!}$

(b) $\sum_{k=0}^{n} k \cdot k! = (n+1)! - 1$

EXERCISE 3.5.9 Define a sequence recursively by letting $a_0 = 0$, and defining $a_{n+1} = \frac{1}{2}a_n + 1$ for $n \geq 0$. Then $a_n = \dfrac{2^n - 1}{2^{n-1}}$ for all $n \geq 0$.

EXERCISE 3.5.10 By Theorem 3.4.16, polynomial multiplication will reveal that

$$(1 - x)(1 + x + x^2 + x^3 + x^4) = 1 - x^5$$

[14] You can prove these without induction if you exploit previous exercises appropriately.

[15] In the inductive step, separate the cases $a = 0$ and $a > 0$. Look to Exercise 2.2.5(j).

[16] Use the fact that $0 \leq -b < -a$ and Exercise 3.5.6.

[17] Use Exercise 2.2.6(c).

If $x \neq 1$, we may divide both sides through by $1 - x$ to have

$$1 + x + x^2 + x^3 + x^4 = \frac{1 - x^5}{1 - x}$$

This algebraic observation suggests a general formula for the sum of powers of a real number $x \neq 1$. If $n \geq 0$, then

$$1 + x + x^2 + x^3 + \cdots + x^n = \sum_{k=0}^{n} x^k = \frac{1 - x^{n+1}}{1 - x} \tag{3.51}$$

Prove Eq. (3.51) with an induction argument on $n \geq 0$.

Exercise 3.5.11 Prove the following factorization formula for nonzero real numbers a and b and nonnegative integer n.[18]

$$a^{n+1} - b^{n+1} = (a - b)(a^n + a^{n-1}b + a^{n-2}b^2 + \cdots + a^2b^{n-2} + ab^{n-1} + b^n) \tag{3.52}$$

Another reason Theorem 3.5.1 is useful is that a result might not be true for all positive integers, but only eventually true, that is, true for $n \geq j$ for some positive integer j. Just like we use $\sum_{k=1}^{n} a_k$ to denote the sum of the a_k, we can use the notation $\Pi_{k=1}^{n} a_k$ to denote the product of the terms. In other words,

$$\Pi_{k=1}^{n} a_k = a_1 \cdot a_2 \cdot a_3 \cdots a_n$$

Exercise 3.5.12 For all $n \geq 2$, $\Pi_{k=2}^{n}\left(1 - \frac{1}{k}\right) = \frac{1}{n}$.

Exercise 3.5.13 If $n \geq 4$ is an integer, then $2^n < n!$.

Exercise 3.5.14 If $n \geq 2$ is an integer, then $3 \mid (n^3 - n)$.

The next lemma does not require induction, but is helpful in the exercise that follows.

Lemma 3.5.15 If $n \geq 3$ is an integer, then $n^2 > 2n + 1$.

Proof. Let $n \geq 3$ be an integer. Then

$$n^2 = n \cdot n \geq 3n = 2n + n > 2n + 1$$

□

Exercise 3.5.16 If $n \geq 5$ is an integer, then $n^2 < 2^n$.

[18] Don't make another induction argument. Letting $x = b/a$ in Eq. (3.51) provides a good start.

Strong Induction

There is another way to build the set S on p. 80. Consider the following. Suppose S is a subset of the positive integers that has the following properties.

(K1) $1 \in S$

(K2) If $n \geq 2$ and $1, 2, \ldots, n-1 \in S$, then $n \in S$.

Then by yet another mimicking of the proof of Theorem 3.4.6, we could show that $S = \mathbb{N}$. This is called the *strong principle of mathematical induction*, or SPMI.

Theorem 3.5.17 (SPMI). Suppose S is a subset of the positive integers that has properties K1–K2. Then $S = \mathbb{N}$.

Exercise 3.5.18 This exercise addresses why Theorem 3.5.17 is called strong induction. To do so, we alter property I2 slightly to read: If $n \geq 2$ and $n-1 \in S$, then $n \in S$. For each of the following pairs of statements, state which is stronger.

(a)	$n \geq 2$ and $n-1 \in S$	$n \geq 2$ and $1, 2, \ldots, n-1 \in S$
(b)	Property I2	Property K2
(c)	Properties I1–I2	Properties K1–K2
(d)	If I1 and I2, then $S = \mathbb{N}$.	If K1 and K2, then $S = \mathbb{N}$.

This is why the SPMI can be more powerful than the PMI. If we are required to prove a result is true for all $n \geq 1$ and induction seems to be the way to go, we might find that regular induction does not provide us with a strong enough assumption to make the inductive leap. With either form of induction, we would still need to show that $1 \in S$. But to make the inductive step, regular induction would only allow us to assume $n \in S$ and require us to conclude $n+1 \in S$ solely from this. On the other hand, strong induction allows us to assume that all of $1, 2, \ldots, n-1 \in S$, and requires us to conclude $n \in S$ from this more extensive set of assumptions.[19]

In Section 2.5 we defined a positive integer to be prime provided it has exactly two distinct positive integer factors. An integer $n \geq 2$ that is not prime is called *composite*, and such a number can be written in the form $n = ab$, where a and b are positive integers strictly less than n. The following theorem addresses the factorization of positive integers into a product of primes. We use a slight variation of the SPMI by rooting it at $n = 2$.

Theorem 3.5.19 (Fundamental Theorem of Arithmetic). Every integer $n \geq 2$ can be written as the product of primes, and this factorization is unique, except perhaps for the order in which the factors are written.

[19] Actually, strong induction is no stronger than regular induction. They are logically equivalent.

Proof. We use strong induction on $n \geq 2$ to show existence. Since 2 is prime, the result is true for $n = 2$. So let $n \geq 3$ be given, and suppose all of $2, 3, \ldots, n-1$ can be written as the product of primes. Now if n is itself prime, then its prime factorization is trivial, and the result is true. On the other hand, if n is composite, then there exist positive integers a and b such that $1 < a < n$, $1 < b < n$, and $n = ab$. By the inductive assumption, since $2 \leq a, b \leq n-1$, both a and b can be written as the product of primes. That is

$$a = p_1 p_2 \cdots p_s$$

$$b = q_1 q_2 \cdots q_t$$

where all p_k and q_k are prime. Thus $n = p_1 \cdots p_s q_1 \cdots q_t$, and we have found a prime factorization for n.

To show uniqueness, choose $n \geq 2$, and suppose $n = p_1 p_2 \cdots p_k$ and $n = q_1 q_2 \cdots q_l$ are two ways of writing n as a product of primes. We show by induction on k that $l = k$ and, with some possible reordering of the q_i, that $p_i = q_i$ for all $1 \leq i \leq k$.

If $k = 1$, then $n = p_1$, so that n is prime. Thus $l = 1$ and $p_1 = q_1$. So suppose $k \geq 2$ and suppose that any factorization of a positive integer into $k-1$ primes is unique up to order of the factors. Since

$$p_1 p_2 \cdots p_k = q_1 q_2 \cdots q_l \tag{3.53}$$

then $p_k \mid q_1 q_2 \cdots q_l$. By Exercise 3.4.19, there exists some j, where $1 \leq j \leq l$, such that $p_k \mid q_j$. Since q_j is prime, $p_k = q_j$. Reordering the right-hand side of Eq. (3.53) by switching q_l and q_j, then canceling p_k and q_l, we have factorizations of the positive integer n/p_k into $k-1$ and $l-1$ primes. Since $p_k \geq 2$, we have that $n/p_k < n$. Applying the inductive assumption to n/p_k, we conclude that $k - 1 = l - 1$, and with some possible reordering of $q_1, \ldots, q_{l-1}$, we have that $p_i = q_i$ for $1 \leq i \leq l-1$. Thus $l = k$, the factorization of n into k primes is unique up to the order of the factors. □

Strong induction on n can be adapted slightly in the following way. Establish a range of base cases $0 \leq n \leq k-1$, then use a strong inductive assumption for $n \geq k$.

EXERCISE 3.5.20 Every integer $n \geq 30$ can be written as an integer linear combination of 6 and 7 using nonnegative multiples. That is, if $n \geq 30$ is an integer, then there exist nonnegative integers j and k such that $n = 6j + 7k$.[20]

EXERCISE 3.5.21 Let a and b be integers, where $a > 0$ and $b \geq 0$. Prove the existence part of the division algorithm (Theorem 2.4.7) using an adapted strong induction on b with base cases $0 \leq b \leq a-1$.

[20] Establish base cases from $n = 30$ to $n = 35$.

EXERCISE 3.5.22 Suppose $n \geq 1$ disks that are black on one side and white on the other are arranged in a straight line with a random arrangement of black sides up. A game is to be played whereby a black disk is removed, and its immediate neighbors, if any, are flipped over. (Two disks are not said to be immediate neighbors if there have been any empty spaces created between them.) A game is winnable if it is possible to remove all n disks. What constitutes a winnable game? Prove your claim using strong induction.[21]

EXERCISE 3.5.23 Repeat Exercise 3.5.22 for a similar game where $n \geq 3$ disks are arranged in a circle.

3.6 Equivalence Relations

The notion of equality is so central to the mathematical enterprise that we do not even think about the ways we manipulate equations. In Chapter 0 we stated three assumptions about equality in the real numbers.

(A1) **Properties of Equality:**

(a) For all $a \in \mathbb{R}, a = a$. (Reflexive)

(b) For all $a, b \in \mathbb{R}$, if $a = b$, then $b = a$. (Symmetric)

(c) For all $a, b, c \in \mathbb{R}$, if $a = b$ and $b = c$, then $a = c$. (Transitive)

In this section, we will consider a variety of sets and define on them a notion of *equivalence* that has the reflexive, symmetric, and transitive properties. As you work through these examples and exercises, similarities will begin to show through, and these will illustrate other abstract features of equivalences that we will explore in Section 3.7.

Before we get to the examples, we note that the symbol for equivalence can vary greatly from situation to situation. Given a set S and two elements $x, y \in S$, there are different symbols commonly used to denote that x and y are equivalent, whatever that might mean in its particular context. For example,

$$x = y$$
$$x \equiv y$$
$$x \equiv_n y$$
$$x \sim y$$

[21] David Beckwith, Glenn G. Chappell, and O. P. Lossers, *American Mathematical Monthly* Vol. 104, 9 (Nov. 1997): 876.

$$x \cong y$$

$$x \leftrightarrow y$$

$$x \doteq y$$

$$x \Leftrightarrow y$$

$$x\mathrm{R}y$$

are but a few. The statement $x \mathrm{R} y$ is read "x is related to y," and derives from a way of addressing equivalence in terms of a *relation*, which we will investigate in Section 3.11. Since our goal is to address the features of equivalence that transcend context, let's choose one of these common symbols, say $\equiv$, and stick with it.

Definition 3.6.1 Let S be a non-empty set, and let "$x \equiv y$" be a statement for all $x, y \in S$. That is, for every $x, y \in S$, either $x \equiv y$ or $x \not\equiv y$ is true. Suppose also that $\equiv$ has the following properties on S.

(E1) For all $x \in S, x \equiv x$. (Reflexive)

(E2) For all $x, y \in S$, if $x \equiv y$, then $y \equiv x$. (Symmetric)

(E3) For all $x, y, z \in S$, if $x \equiv y$ and $y \equiv z$, then $x \equiv z$. (Transitive)

Then $\equiv$ is called an *equivalence relation* on S.

Here is our first concrete example. On the surface, it might not sound particularly mathematical. To the contrary, it is, shall we say, equivalent to all the other examples we will study, in that it has properties E1–E3.

Example 3.6.2 Let C be the set of all cities on earth. Suppose for the sake of argument that there are no one-way roads. Define two cities $x, y \in C$ to be equivalent, $x \equiv y$, provided it is possible to drive on roads *from* city x *to* city y. Then $\equiv$ defines an equivalence relation on C.

Solution We must verify that properties E1–E3 hold on C.

(E1) Pick $x \in C$. Since it is possible (trivially) to drive from city x to city x on roads, then $x \equiv x$.

(E2) Pick $x, y \in C$ and suppose $x \equiv y$. Then it is possible to drive from x to y on roads. Since we have assumed no roads are one-way, then it is also possible to drive from y to x, since the same roads going from x to y can be traveled in the opposite direction. Thus $y \equiv x$.

(E3) Pick $x, y, z \in C$ and suppose $x \equiv y$ and $y \equiv z$. Then it is possible to drive from x to y on roads, and it is possible to drive from y to z on roads. By beginning at x, driving to y, then to z, we see that it is possible to drive from x to z on roads. Thus $x \equiv z$.

Since $\equiv$ satisfies properties E1–E3, we have that $\equiv$ is an equivalence relation on C. ■

EXERCISE 3.6.3 Name three cities that are equivalent (according to Example 3.6.2) to Beira.

Example 3.6.2 and those to follow are all cut from the same pattern. They involve first defining a criterion by which two elements of a set are declared to be equivalent, then demonstrating that this relationship satisfies properties E1–E3. To illustrate the general pattern, let's use the expression $P(x, y)$ to represent a generic statement involving x and y. Notice that the two indeterminates x and y are ordered in $P(x, y)$.

Definition 3.6.4 (Sample) For $x, y \in S$, define $x \equiv y$ provided $P(x, y)$.

Theorem 3.6.5 (Sample). The equivalence $\equiv$ in Definition 3.6.4 is an equivalence relation on S.

Proof. We show that properties E1–E3 hold for $\equiv$ on S.

(E1) Pick $x \in S$. Then ... so that $P(x, x)$. Thus $x \equiv x$.

(E2) Pick $x, y \in S$ and suppose $x \equiv y$. Then $P(x, y)$. Thus ..., so that $P(y, x)$. Therefore, $y \equiv x$.

(E3) Pick $x, y, z \in S$ and suppose $x \equiv y$ and $y \equiv z$. Then $P(x, y)$ and $P(y, z)$. Then ..., so that $P(x, z)$. Thus $x \equiv z$.

Since $\equiv$ satisfies properties E1–E3, we have that $\equiv$ defines an equivalence relation on S. □

Here is one more example in full detail. The expression $\equiv_{\mathbb{Z}}$ is read "equivalence modulo $\mathbb{Z}$."

Example 3.6.6 For real numbers x and y, define $x \equiv_{\mathbb{Z}} y$ provided $x - y$ is an integer. Show that $\equiv_{\mathbb{Z}}$ defines an equivalence relation on the real numbers.

Solution We show that $\equiv_{\mathbb{Z}}$ satisfies properties E1–E3.

(E1) Pick $x \in \mathbb{R}$. Then $x - x = 0$, which is an integer. Thus $x \equiv_{\mathbb{Z}} x$.

(E2) Pick $x, y \in \mathbb{R}$ and suppose $x \equiv_{\mathbb{Z}} y$. Then $x - y$ is an integer. Thus $y - x = -(x - y)$ is also an integer, so that $y \equiv_{\mathbb{Z}} x$.

(E3) Pick $x, y, z \in \mathbb{R}$, and suppose $x \equiv_{\mathbb{Z}} y$ and $y \equiv_{\mathbb{Z}} z$. Then $x - y$ and $y - z$ are both integers. Since the integers are closed under addition, we have that $x - z = (x - y) + (y - z) \in \mathbb{Z}$. Thus $x \equiv_{\mathbb{Z}} z$.

Since $\equiv_{\mathbb{Z}}$ satisfies properties E1–E3, we have that $\equiv_{\mathbb{Z}}$ is an equivalence relation on the real numbers. ■

If some inquisitive soul were to ask for some real numbers that are equivalent to 8.363 according to the definition in equivalence in Example 3.6.6, we could take the following approach. Since $x \equiv_{\mathbb{Z}} y$ means that $x - y$ is an integer, we know that $x \equiv_{\mathbb{Z}} 8.363$ provided $x - 8.363$ is an integer. Now if your three favorite integers are, respectively, 15, −188, and 0, we would have

$$x - 8.363 = 15 \quad x - 8.363 = -188 \quad \text{and} \quad x - 8.363 = 0 \tag{3.54}$$

which would yield 23.363, −179.637, and 8.363. With practice, you might develop a quicker approach of your own.

Exercise 3.6.7 Let S be the xy-plane, that is, $S = \{(x, y) : x, y \in \mathbb{R}\}$. For two points (x_1, y_1) and (x_2, y_2), define the two points to be equivalent, written $(x_1, y_1) \equiv (x_2, y_2)$, provided $y_2 - y_1 = 3(x_2 - x_1)$. Show that $\equiv$ defines an equivalence relation on S. Name three elements of S that are equivalent to $(2, 1)$.

Exercise 3.6.8 Let $\mathcal{F}$ be the family of all sets. Show that Definition 3.1.9 defines an equivalence relation on $\mathcal{F}$.

To show that $\equiv$ is an equivalence relation on a set involves showing that all three properties E1–E3 are satisfied. Furthermore, each of these properties is a universal statement about the elements of the set. Thus to show that $\equiv$ is not an equivalence relation, we only need to show that one of the properties fails, and this demonstration will involve constructing a counterexample.

Example 3.6.9 For points in the xy-plane, define $(x_1, y_1) \equiv (x_2, y_2)$ provided either $x_1 = x_2$ or $y_1 = y_2$. Then $\equiv$ is not an equivalence relation because it is not transitive. For example, $(1, 2) \equiv (3, 2)$ and $(3, 2) \equiv (3, 4)$, but $(1, 2) \not\equiv (3, 4)$. ■

Exercise 3.6.10 For integers m and n, define $m \equiv n$ provided $|m - n| \neq 1$. Is $\equiv$ an equivalence relation on the whole numbers? Prove or disprove.

Exercise 3.6.11 For positive integers m and n, define $m \equiv n$ provided $mn \geq 2$. Is $\equiv$ an equivalence relation? Prove or disprove.

Exercise 3.6.12 Let S be the set of four-letter words in the *Random House Unabridged Dictionary*. For two words $x, y \in S$ define $x \equiv y$ if it is possible to construct a sequence of words in S, beginning with x and changing one letter at a time to create another word in S at each step, ending up at y. (A sequence of one word is considered valid.) For example BABY $\equiv$ POOP because of the sequence BABY, BABE, BARE, BARN, BURN, BURP, BUMP, PUMP, POMP, POOP. Is $\equiv$ an equivalence relation on S? Prove or disprove.

Exercise 3.6.13 Let $\mathcal{F}$ be the family of all sets. For sets A and B, define $A \equiv B$ provided either $A \subseteq B$ or $B \subseteq A$. Is $\equiv$ an equivalence relation on $\mathcal{F}$? Prove or disprove.

Exercise 3.6.14 Suppose we have a square that sits in the xy-plane with its center at the origin and rotated so that its corners lie on the coordinate axes. Suppose also that we have written the numbers $\{1, 2, 3, 4\}$, one on each corner, as in Figure 3.10. (Assume the numbers are visible on the front and back of the square. All we care about is which number is in each of the four possible positions.) Call one such writing of these numbers an *assignment*. For two assignments x and y, define $x \equiv y$ if it is possible to convert x into y by rotating the square around its center point and/or flipping the square over about the x-axis. Does $\equiv$ define an equivalence relation on the set of all assignments? Prove or disprove.

Exercise 3.6.15 Let $\mathcal{F}$ be the family of all non-empty sets. For two sets $A, B \in \mathcal{F}$, define $A \equiv B$ provided $A \cap B$ is non-empty. Is $\equiv$ an equivalence relation on $\mathcal{F}$? Prove or disprove.

Exercise 3.6.16 Is $\leq$ an equivalence relation on the real numbers?

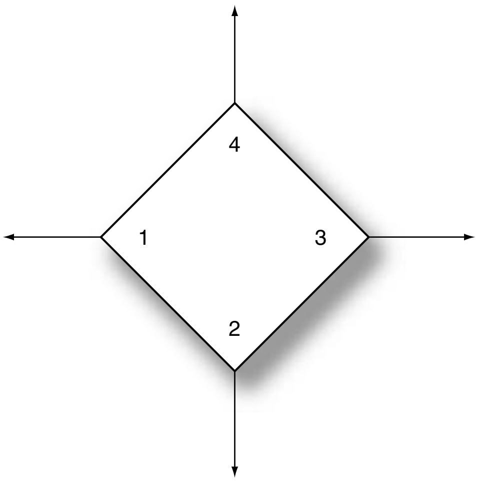

Figure 3.10 An assignment of numbers to the square.

EXERCISE 3.6.17 Is $\neq$ an equivalence relation on the real numbers?

Our next example of an equivalence relation is very important in abstract algebra, and it pervades Part III of this text. The symbol $\equiv_6$ is read "equivalence modulo 6" (or mod 6).

EXERCISE 3.6.18 For integers x and y, define $x \equiv_6 y$ provided there exists an integer k such that $x - y = 6k$. This definition says that $x \equiv_6 y$ if $x - y$ is a multiple of 6. Then $\equiv_6$ defines an equivalence relation on the integers.

EXERCISE 3.6.19 Construct the set consisting of all integers that are equivalent (modulo 6) to 14.

EXERCISE 3.6.20 What is the smallest nonnegative integer that is equivalent (mod 6) to -16?

The number 6 in our definition of equivalence (mod 6) was chosen arbitrarily, of course. We can discuss $\equiv_n$ for any positive integer n, where we define $x \equiv_n y$ provided there exists an integer k such that $x - y = nk$. Notationally there are two other common ways of writing this form of equivalence among integers. They are

$$x \equiv y \pmod{n} \tag{3.55}$$

$$x \equiv y \quad (n) \tag{3.56}$$

EXERCISE 3.6.21 What is the smallest nonnegative integer x such that $1195 \equiv_{11} x$?

EXERCISE 3.6.22 Suppose x is an integer such that $x \equiv_9 65$. What is the smallest nonnegative integer y such that $x \equiv_9 y$?

Equivalence mod n has several important features that we investigate here. First, notice that $81 \equiv_7 25$ because $81 - 25 = 56 = 8 \times 7$. Applying the division algorithm to $a = 7$ and $b_1 = 81$, then to $a = 7$ and $b_2 = 25$, we get $81 = 7 \times 11 + 4$ and $25 = 7 \times 3 + 4$. So 81 and 25 have the same remainder when divided by 7 according to the division algorithm. This illustrates the following theorem, which says merely that x and y differ by a multiple of n if and only if they have the same remainder when divided by n.

EXERCISE 3.6.23 Suppose x and y are integers and n is a positive integer. Then $x \equiv_n y$ if and only if x and y have the same remainder when divided by n according to the division algorithm.

EXERCISE 3.6.24 Let n be a given positive integer. Then for every integer x, there exists $y \in \{0, 1, \ldots, n - 1\}$ such that $x \equiv_n y$.

EXERCISE 3.6.25 Suppose x and y are integers such that $x \equiv_3 y \not\equiv_3 0$. Then $xy \equiv_3 1$.[22]

Sometimes a mathematician doing research believes a certain theorem is true and can be proved, but he is mistaken. Often, he works his way through a proof only to find a hole in his reasoning, a gap that he cannot cross to arrive at the desired conclusion. Analyzing this gap might lead to the discovery of a counterexample to his supposed theorem, which in turn might also lead to a somewhat weaker theorem that *can* be proved for all cases except for those that are like the counterexample.

In a similar way, textbook authors sometimes think they are devising a really cool exercise, only to discover a mistake that they would be embarrassed for their colleagues to know about. The last exercises of this section illustrate such a mistake that an unnamed author[23] made, a principle discovered that generates counterexamples, and a subsequent weaker theorem that can be proved.

EXERCISE 3.6.26 Let $S = \{(x, y) : x, y \in \mathbb{Z}\}$, and define $(x_1, y_1) \equiv (x_2, y_2)$ provided $x_1 y_2 = x_2 y_1$. Try to show that $\equiv$ defines an equivalence relation on S. Where does the proof break down? Can you use the obstacle to discover a counterexample to the claim? If not, look to the next exercise.

EXERCISE 3.6.27 Show that transitivity fails in the definition of equivalence in Exercise 3.6.26.[24]

EXERCISE 3.6.28 The set S in Exercise 3.6.26 can be redefined to eliminate the one problematic element noted in Exercise 3.6.27, so that the definition of equivalence will satisfy properties E1–E3. Redefine S appropriately, and then use your work from Exercise 3.6.26 and your new definition of S to prove that $\equiv$ is an equivalence relation on your newly defined set.

3.7 Equivalence Classes and Partitions

In Example 3.6.2, it might have occurred to you that defining two cities to be equivalent in terms of accessibility by roads takes all cities on earth and lumps them together into groups of cities that are mutually accessible from each other. Presumably, for example, there is a network of two-way roads connecting all cities on the north island of New Zealand. Thus it is possible to drive from any city on the island to any other. As a result, every city on the north island is equivalent to

[22] If x and y are not equivalent to 0 (mod 3), then what possibilities remain?

[23] I myself.

[24] There is one point that every element of S is equivalent to, but there do exist other points that are not equivalent to each other.

every city on the north island, including itself. Furthermore, no city outside the north island of New Zealand is equivalent to any city on that island. If that were to change, say by some newly constructed bridge between Auckland, New Zealand, and Santiago, Chile, then Santiago would become equivalent to all the cities on the north island of New Zealand. Furthermore, all cities accessible from Santiago would also become equivalent to all cities on the north island of New Zealand. It is sort of like the molecules in two drops of water. Either the two drops do not touch at all, or, if they do, they instantly merge into one drop.

This illustrates that there is a lot of strength in the properties E1–E3. One very important feat that an equivalence relation performs is that it completely splits the set up into non-empty, nonoverlapping subsets. In general, the splitting of a set into nonoverlapping subsets is called *partitioning*. When an equivalence relation is defined on a set, it naturally partitions the set into subsets where elements of the same subset are all equivalent to each other, and any two elements of different subsets are not equivalent to each other.

In this section, we will define partitions in general and construct some examples. Then we will see how an equivalence relation gives rise to a partition of the set on which it is defined.

Definition 3.7.1 Suppose S is a set and $\mathcal{F} = \{A\}$ is a family of subsets of S. Then $\mathcal{F}$ is said to be a *partition* of S if

(P1) Every A in $\mathcal{F}$ is non-empty;

(P2) If $A, B \in \mathcal{F}$ and $A \cap B$ is non-empty, then $A = B$; and

(P3) $\bigcup_{A\in\mathcal{F}} A = S$.

Let's illustrate with a specific example before we discuss some general properties of partitions. Just remember, a partition of a given set S is nothing but a family of subsets of S with some special properties.

Example 3.7.2 Consider the following four families of subsets of $\mathbb{N}_{10}$.

$$\begin{aligned}
\mathcal{F}_1 &= \big\{\{1\}, \{2, 3, 5, 7\}, \{4, 6, 8, 10\}, \{9\}\big\} \\
\mathcal{F}_2 &= \big\{\emptyset, \{1, 2, 3, 4\}, \{5, 6, 7, 8\}, \{9, 10\}\big\} \\
\mathcal{F}_3 &= \big\{\{1\}, \{2, 4, 6, 8, 10\}, \{3, 6, 9\}, \{5, 10\}, \{7\}\big\} \\
\mathcal{F}_4 &= \big\{\{\text{Primes in } \mathbb{N}_{10}\}, \{\text{Composites in } \mathbb{N}_{10}\}\big\}
\end{aligned} \tag{3.57}$$

Of these four families, only $\mathcal{F}_1$ is a partition of $\mathbb{N}_{10}$. Notice how $\mathcal{F}_1$ satisfies properties P1–P3. First, no set in $\mathcal{F}_1$ is empty. Second, they are all disjoint. If we choose any two of them and they have a non-empty intersection, then they are the same set. Third, forming the union across $\mathcal{F}_1$ produces $\mathbb{N}_{10}$.

Notice how the other families fail to be a partition of $\mathbb{N}_{10}$. Since a partition of a set must contain only non-empty subsets, $\mathcal{F}_2$ fails to be a partition of $\mathbb{N}_{10}$. In $\mathcal{F}_3$, there exist two distinct sets with non-empty intersection. Finally, the union across $\mathcal{F}_4$ is not $\mathbb{N}_{10}$ for 1 is neither prime nor composite. ■

In the work we will do in this section, there will be some program by which we construct a family of subsets of a given set in order to create a partition. To verify that this family does indeed partition the set, we must show that the family has properties P1–P3.

(P1) The program for forming the subsets of S must always generate non-empty sets. Thus if we choose an arbitrary set in the family, we must be able to find some element of S in the chosen set.

(P2) No element of S can be in more than one distinct set in the family. Thus if we choose sets A and B in the family and suppose they have a common element ($A \cap B \neq \emptyset$), we must be able to show $A = B$.

(P3) This property says that the generated sets in the family completely exhaust all elements of S. Naturally, the program for forming the family of subsets of S should be defined so that it creates only subsets of S. Then Exercise 3.3.13 will guarantee that the $\subseteq$ property P3 is satisfied. What remains to be shown is $\supseteq$, which amounts to showing that every element of S is in some set in the family.

Example 3.7.3 From Example 3.6.2, create a subset of C in the following way. Choose a city x, and define A_x to be the set of all cities on earth from which there are roads to city x. Now let

$$\mathcal{F} = \{A_x : x \in C\} \tag{3.58}$$

the family of all these subsets of C constructed by letting x be every city on earth. Then $\mathcal{F}$ is a partition of C.

Solution We verify that properties P1–P3 are satisfied on $\mathcal{F}$.

(P1) Pick any $A \in \mathcal{F}$. By the way sets in $\mathcal{F}$ are constructed, there exists $x \in C$ for which $A = A_x$. Since city x has roads to itself, we have that $x \in A$, so that A is non-empty.

(P2) Pick $A_x, A_y \in \mathcal{F}$, and suppose $A_x \cap A_y \neq \emptyset$. Then there exists $z \in A_x \cap A_y$. We show that $A_x = A_y$. First pick $a \in A_x$. Then it is possible to drive from city a to city x on roads. Now since $z \in A_x$, there are roads from z to x, so that it is also possible to drive from x to z. (Draw a picture!) Furthermore, since $z \in A_y$, there are roads from z to y. Thus we may travel $a \to x \to z \to y$, and have that it is possible to drive from a to y. Thus $a \in A_y$, so that $A_x \subseteq A_y$. By an identical argument in the reverse direction $A_y \subseteq A_x$. Thus $A_x = A_y$.

(P3) Since $A_x \subseteq C$ for all $x \in C$, Exercise 3.3.13 implies that $\cup_{x \in C} A_x \subseteq C$. To show $\supseteq$, pick any $y \in C$. Now $A_y \in \mathcal{F}$ and $y \in A_y$. Thus $y \in \cup_{x \in C} A_x$, and $C \subseteq \cup_{x \in C} A_x$. ■

Among the islands along the inside passage of southeast Alaska is Wrangell Island, whose only town, Wrangell, has a population of around 2500. Though there are roads between Wrangell and remote parts of the island, it is not possible to get to Wrangell on roads from any other city on earth. Thus there is only one element of the set A_{Wrangell}, and that is Wrangell. In the construction of $\mathcal{F}$ in Eq. (3.58), we allowed x to take on the names of all cities on earth to generate all the A_x in $\mathcal{F}$. Of the vast number of cities on earth, each of which generated its own subset of C, A_{Wrangell} was generated only once, when $x = \text{Wrangell}$.

Contrast this with A_{Dallas}. Dallas is an element of A_x for thousands of $x \in C$ because there are so many cities accessible by roads from Dallas. So in the construction of $\mathcal{F}$, Dallas is placed into such sets as A_{Lubbock}, A_{Banff}, and $A_{\text{Amsterdam}}$.[25] However, in spite of this redundancy, $\mathcal{F}$ partitions C because all the sets containing Dallas are the same.

EXERCISE 3.7.4 In Example 3.6.6 we defined an equivalence relation on the real numbers by defining $x \equiv_{\mathbb{Z}} y$, provided $x - y$ is an integer. For a given real number x, define the following subset of the real numbers:

$$A_x = \{y \in \mathbb{R} : y \equiv_{\mathbb{Z}} x\} = \{y \in \mathbb{R} : y - x \in \mathbb{Z}\} \tag{3.59}$$

Create a family of subsets of real numbers by letting $\mathcal{F} = \{A_x : x \in \mathbb{R}\}$.

(a) Describe $A_{4.34}$ by listing some of its elements.

(b) Show that $\mathcal{F}$ partitions the real numbers by showing it has properties P1–P3.

Now let's see how an arbitrary equivalence relation on a non-empty set S is tied to a partition of S. We will work our way into a theorem tying them together by way of some examples. First, a definition that presents standard mathematical notation for the sets we have thus far written as A_x.

Definition 3.7.5 Suppose $\equiv$ defines an equivalence relation on a non-empty set S. For an element $x \in S$, define

$$[x] = \{\, y \in S : y \equiv x \,\}$$

That is, $[x]$ contains all elements of S that are equivalent to x. This subset of S is called the *equivalence class* of x.

---

[25] Amsterdam, Missouri, population 281.

Example 3.7.6 From Example 3.6.2, the equivalence class of Auckland is the set of all cities with roads leading to Auckland, that is, all cities on the north island of New Zealand. It is the same as the equivalence class of Wellington. ■

Example 3.7.7 For equivalence mod 6 from Exercise 3.6.18, given any integer x,

$$[x] = \{y \in \mathbb{Z} : y - x = 6k, \quad \text{for some } k \in \mathbb{Z}\}$$

Thus $[x]$ is the set of all integers y of the form $y = 6k + x$, where k takes on all possible integer values. ■

EXERCISE 3.7.8 From Example 3.7.7, describe $[-4]$ by listing some of its elements. How many distinct equivalence classes are there in this example? Describe all of them by listing some elements from each equivalence class.

EXERCISE 3.7.9 From Example 3.6.14, how many assignments are in the equivalence class of the assignment in Figure 3.10?

In all the preceding examples of equivalence relations, the equivalence classes seem to be a basis of a partition of the set. This is no coincidence, of course. Theorem 3.7.10 says that properties E1–E3 can be used to demonstrate that the family of equivalence classes satisfies P1–P3. Most details are left to you as an exercise.

Theorem 3.7.10 Suppose $\equiv$ defines an equivalence relation on a set S. Define $\mathcal{F} = \{\,[x] : x \in S\,\}$. That is, $\mathcal{F}$ is the collection of all equivalence classes generated by letting x take on all values in S. Then $\mathcal{F}$ is a partition of S.

Proof. We verify that properties P1–P3 from Definition 3.7.1 are satisfied.

(P1) Pick any $[x] \in \mathcal{F}$. Since $x \equiv x$, it follows that $x \in [x]$. Thus $[x]$ is not empty.

(P2) Pick $[x], [y] \in \mathcal{F}$ and suppose $[x] \cap [y]$ is not empty . . . .

(P3) Since every set in $\mathcal{F}$ is defined to be a subset of S, we have $\cup_{x \in S}[x] \subseteq S$ by Exercise 3.3.13. So we must show $\supseteq$. Pick $y \in S$ . . . . □

EXERCISE 3.7.11 Complete parts P1–P3 of the proof of Theorem 3.7.10.

Theorem 3.7.10 is very important for the following reason. If some form of equivalence is defined on a set, and it can be verified that this definition is an equivalence relation, then Theorem 3.7.10 guarantees that the set is automatically partitioned into equivalence classes. We can then think of any element of a particular equivalence class as being representative of all elements in its class. Furthermore, we can think of this equivalence in precisely the same way we think of equality, where two equivalent elements are, in many senses, interchangeable.

The partitioning of a set into equivalence classes lumps the elements of the set into categories where one element can be replaced with any other element in its category, within limits, of course. In Section 3.8 we will illustrate this phenomenon on the rational numbers, where we investigate the familiar definition of equality of two fractions. Then we will show how the binary operations of addition and multiplication of rational numbers are defined so that different representative elements of equivalence classes are indeed interchangeable.

3.8 Building the Rational Numbers

As an illustration of the concept of equivalence relation, we want to take a look at equality, addition, and multiplication in the integers and see how they give rise to equality, addition, and multiplication in the rationals. This section effectively dissects of some of the real number properties assumed in Chapter 0 and proved in Chapter 2. Thus a few prefatory words are in order about what we are trying to accomplish here.

First, consider the whole numbers. Beginning with the set $\{0, 1\}$, which by assumption A15 is indeed a two-element set, we calculate all possible sums of its elements. We know from the behavior of 0 that $0 + 0 = 0$ and $1 + 0 = 0 + 1 = 1$. What is not immediately clear is $1 + 1$. Is it possible that $1 + 1 = 0$ or $1 + 1 = 1$? Certain assumptions and properties of real numbers allow you to prove that neither of these is possible.

Exercise 3.8.1 Show that $1 + 1 = 0$ and $1 + 1 = 1$ are impossible.[26]

Since $1 + 1$ is neither 0 nor 1, it might be more efficient to create a new symbol, like 2 perhaps, to denote $1 + 1$. This is just the beginning of a process of building the whole numbers as a set of *successors* of previously defined whole numbers. The assumptions used to build them are called *Peano's postulates*. Because they are formulated without reference to the set of real numbers as a context, Peano's postulates and their implications generate the set of whole numbers in a somewhat more complicated way than we describe here. One of the assumptions is that zero is not a successor to any whole number, so that the sequence of successors, which we call $1, 2, 3, \ldots$, never circles back around to zero.

Once the set of whole numbers has been constructed, we extend it to the integers by tossing in the additive inverses of all the whole numbers. These additive inverses are indeed not whole numbers themselves, except for -0. The reason is that if any positive whole number a had a whole number b as its additive inverse, then $a + b = 0$ would imply that 0 is the successor to $a + b - 1$.

In this section, we want to give ourselves the standard assumptions about equality and binary operations as applied only to the integers, as we did in Chapter 0

[26] Consider the trichotomy law and cancellation of addition.

for the real numbers. Specifically, let's make the following assumptions about the integers, stated in a form very condensed from that in Chapter 0.

(Z1) Integer equality is an equivalence relation and therefore has properties E1–E3.

(Z2) Integer addition and multiplication are well defined, closed, associative, and commutative binary operations. Furthermore, the integers contain the additive and multiplicative identities, the integers are closed under additive inverses, and multiplication distributes over addition.

(Z3) The positive integers are closed under addition and multiplication.

(Z4) All relevant results from Chapter 2, as derived from the previous assumptions, apply to the integers.

In what follows, we want to create the rational numbers from scratch, using the integers as building blocks.[27] First we construct the set of rationals; then we define a notion of equivalence on them using integer equality as our criterion, and this definition will coincide with the way we already think of rational equality. Then we will define addition and multiplication on the rationals and show that they are well-defined operations. Throughout this discussion, we will use the term *equality* exclusively with reference to integers, and *equivalence* when we refer to the similar notion we are defining on the rational numbers.

3.8.1 Defining Rational Equality

In Chapter 0, we defined the set of rational numbers as

$$\mathbb{Q} = \{\, p/q : p, q \in \mathbb{Z}, q \neq 0 \,\} \tag{3.60}$$

We normally think of the expression p/q as meaning "p divided by q." For now, we want to dispense with this meaning and conceive of rational numbers as merely ordered pairs of integers, where the second integer is not allowed to be zero. In this discussion, the forward slash in p/q is intended to have no meaning, except that it separates the first integer from the second.

There is nothing inherent in the definition of the rationals in Eq. (3.60) that suggests how you would address the equivalence of two of its elements. In fact, there is more than one useful way to define equivalence on the rational numbers, but we will look only at the familiar form that coincides with our notion of rational equality. Note that in Definition 3.8.2, we refer to integer equality as =, and we

[27] In a more abstract algebraic setting, this would be called constructing the field of quotients of an integral domain.

use juxtaposition to represent integer multiplication. The equivalence denoted $\equiv$ is the symbol we are defining on the rationals.

Definition 3.8.2 For two rational numbers p/q and r/s, we define $p/q \equiv r/s$ provided $ps = qr$.

Notice the positions of p, q, r, s in Definition 3.8.2. These positions must be retained when we translate between $p/q \equiv r/s$ and $ps = qr$.

Exercise 3.8.3 Answer the following questions as they relate to Definition 3.8.2. All alphabetic symbols represent integers.

(a) The statement $a/m \equiv b/n$ corresponds to what integer equation?

(b) The statement $bc = de$ translates into what rational equivalence statement?

(c) Is $6/(-15) \equiv (-4)/10$? Why or why not?

Exercise 3.8.4 Assuming properties Z1–Z4 on the integers, show that Definition 3.8.2 defines an equivalence relation on the rational numbers.

Thanks to Exercise 3.8.4, it follows from Theorem 3.7.10 that the rational numbers are partitioned into equivalence classes where p/q and r/s are in the same equivalence class provided $ps = qr$. You already have an idea of what these equivalence classes look like.

Exercise 3.8.5 Prove the following about the equivalence classes of rational numbers.

(a) For any integer r and nonzero integers q and s, $0/q \equiv r/s$ if and only if $r = 0$.

(b) For any rational number p/q and any nonzero integer s, $p/q \equiv ps/qs$.

By Exercise 3.8.5, all rational numbers of the form $0/q$ are equivalent to each other. Furthermore, if a/b is rational, a is nonzero, and a and b have a common integer factor s, then we may write $a = ps$ and $b = qs$ to have that $a/b \equiv p/q$. This corresponds to the familiar notion that we may cancel out a common factor of a numerator and denominator and have an equivalent rational number. We may then address an equivalence class of rational numbers by referring to an expression of the form p/q, where p and q are assumed to be relatively prime.

3.8.2 Rational Addition and Multiplication

Now that we have an equivalence relation defined on the rational numbers, we can define a form of addition of rational numbers. We will denote this by $\oplus$ to distinguish it from the integer addition we use to define it.

Definition 3.8.6 Given two rational numbers p/q and r/s, we define addition $\oplus$ in the following way.

$$p/q \oplus r/s = (ps + qr)/qs \tag{3.61}$$

Observe that Eq. (3.61) guarantees that $\oplus$ is closed on the rationals. Since integer addition and multiplication are closed, then $ps + qr$ is also an integer, and the fact that q and s are nonzero integers implies that qs is also a nonzero integer. Thus $(ps + qr)/qs$ is indeed a rational number.

Now we must show that addition in the rational numbers is well defined. Notice in the next exercise how that the equations use rational equivalence ($\equiv$) and addition $\oplus$. When you translate them into their equivalent integer form, you will use standard integer equality, addition, and multiplication.

EXERCISE 3.8.7 Addition of rational numbers as defined in Eq. (3.61) is well defined. That is, if $p/q, r/s, t/u, v/w$ are rational numbers, and if $p/q \equiv r/s$ and $t/u \equiv v/w$, then $p/q \oplus t/u \equiv r/s \oplus v/w$.

EXERCISE 3.8.8 Show that $\oplus$ is commutative.

Notice from Eq. (3.61) that anything from the equivalence class of $0/s$ functions as an additive identity in the rationals, for

$$p/q \oplus 0/s \equiv (ps + q \cdot 0)/qs \equiv ps/qs \equiv p/q \tag{3.62}$$

Also, notice $(-p)/q$ and $p/(-q)$, which are equivalent, function as $-(p/q)$.

Definition 3.8.9 Given two rational numbers p/q and r/s, we define multiplication $\otimes$ in the following way.

$$(p/q) \otimes (r/s) = pr/qs \tag{3.63}$$

Clearly, multiplication is closed and commutative. Proving multiplication is well defined is easier than showing addition is well defined.

EXERCISE 3.8.10 Multiplication of rational numbers as defined in Eq. (3.63) is well defined. That is, if $p/q, r/s, t/u, v/w$ are rational numbers, and if $p/q \equiv r/s$ and $t/u \equiv v/w$, then $(p/q) \otimes (t/u) \equiv (r/s) \otimes (v/w)$.

Notice that any rational number of the form p/p functions as a multiplicative identity, and that all rational numbers of this form are in the same equivalence class. Also, recall that a rational number is an additive identity if and only if it is of the form $0/q$. Thus if p/q is rational and $p \neq 0$, then a multiplicative inverse of p/q exists in the rationals. For $(p/q) \otimes (q/p) \equiv (pq)/(qp) \equiv (pq)/(pq)$. Furthermore, any rational number in the equivalence class of q/p functions as a multiplicative inverse of p/q.

3.9 Roots of Real Numbers

Up to this point, we have made no reference to assumptions A20–A22 from Chapter 0. In this section, we want to explore A22:

(A22) For every positive real number x and any positive integer n, there exists a real solution y to the equation $y^n = x$. Such a solution y is called an nth root of x.

You might wonder why assumption A22 refers to the solution of the equation $y^n = x$ in order to address roots of real numbers instead of using the familiar expression $\sqrt[n]{x}$. To do otherwise is to get the cart before the horse. The expression $\sqrt[n]{x}$ derives its meaning from solutions y to the equation $y^n = x$. Whether such solutions exist, and how many such solutions there are must be derived from A22 and other properties of real numbers that we have proved so far. In some cases, a solution of $y^n = x$ will not exist, and we will not be able to define $\sqrt[n]{x}$ in these cases. In other cases, multiple solutions will exist, so that $\sqrt[n]{x}$ would be ambiguous. Thus we do not want to use the expression $\sqrt[n]{x}$ until we are certain of two things. First, we must be certain that such a real number exists, and second, in case there is more than one real solution to the equation $y^n = x$, we agree on which solution we will always intend. By the end of this section, the expression $\sqrt[n]{x}$ will be clearly and uniquely defined for appropriate values of x and n.

Another point to make here is that assumption A22 is not a standard axiomatic assumption in a rigorous study of the real numbers. It actually follows from the Least Upper Bound axiom. In this text, we accept it without question. Here is the main theorem we want to prove in this section.

Theorem 3.9.1 Let x be any real number and n a positive integer. Then the following hold.

(X1) The equation $y^n = 0$ has the unique solution $y = 0$.

(X2) Concerning even roots:

(a) If $x > 0$, the equation $y^{2n} = x$ has precisely two distinct real number solutions, and these are additive inverses of each other.

(b) If $x < 0$, the equation $y^{2n} = x$ has no real number solution.

(X3) Concerning odd roots, the equation $y^{2n+1} = x$ has a unique solution.

Property X1 should be clear, for certainly $0^n = 0$, and if $y^n = 0$, then $y = 0$ from the principle of zero products. The following exercise will lead you systematically through the proof of claims X2–X3. Remember, A22 states that that for every $x > 0$ and positive integer n, the equation $y^n = x$ has *some* real solution y.

Exercise 3.9.2 Prove Theorem 3.9.1 by demonstrating the following. In all parts, n is a positive integer.

(a) Existence in X2(a): Let $x > 0$, and suppose y_0 is a real number solution to $y^{2n} = x$. (Notice $y_0 \neq 0$.) Show that $-y_0$ is also a solution to $y^{2n} = x$.[28]

(b) Nonexistence in X2(b): Suppose $x < 0$. Explain why there is no real number y such that $y^{2n} = x$.

(c) Existence in X3 for $x < 0$: Suppose $x < 0$. Prove that there exists a real number solution y to $y^{2n+1} = x$.[29]

(d) Nonexistence in X3 of solutions of opposite sign: Suppose y_0 is a solution to $y^{2n+1} = x$. Prove that $x > 0$ if and only if $y_0 > 0$.

(e) Uniqueness in X2(a) and X3 (positive case): Let $x > 0$. If $y_1 > 0$ and $y_2 > 0$ satisfy $y_1^n = x$ and $y_2^n = x$, then $y_1 = y_2$.[30]

(f) Uniqueness in X3 (negative case): Let $x < 0$. If $y_1 < 0$ and $y_2 < 0$ satisfy $y_1^{2n+1} = x$ and $y_2^{2n+1} = x$, then $y_1 = y_2$.[31]

With Theorem 3.9.1, we can now introduce the notation $\sqrt[n]{x}$ and define it unambiguously.

Definition 3.9.3 If x is a real number and n is a positive integer, then $\sqrt[2n+1]{x}$ is defined to be the unique solution y of the equation $y^{2n+1} = x$. If $x \geq 0$, the expression $\sqrt[2n]{x}$ is defined to be the unique, nonnegative solution y of the equation $y^{2n} = x$.

Theorem 3.9.1 and Definition 3.9.3 lead us to the following principles of algebraic manipulation. For $x \geq 0$, $y^{2n} = x$ if and only if $y = \pm \sqrt[2n]{x}$. Similarly, for all real numbers x, $y^{2n+1} = x$ if and only if $y = \sqrt[2n+1]{x}$.

Exercise 3.9.4 Suppose $x, y \geq 0$. Show that if $x < y$, then $\sqrt{x} < \sqrt{y}$.[32]

3.10 Irrational Numbers

Accompanied by several axioms we have not discussed in this text, the set $\{0, 1\}$ gives rise to the whole numbers, which is the smallest superset of $\{0, 1\}$ that is closed under addition and multiplication. However, the whole numbers are not closed under additive inverses, so we can expand them to include these additive

[28] Exercise 3.5.6 should come in handy.

[29] Use the fact that $-x > 0$.

[30] Use the factorization formula in Exercise 3.5.11.

[31] Apply part (e) to $-y_1$ and $-y_2$.

[32] Exercise 2.2.6(a) should make this quick.

inverses, arriving at the integers. The integers are not closed under multiplicative inverses, so with a program like that in Section 3.8, we extend the integers to the rationals, as they were defined in Eq. (3.60).

The extension of the whole numbers to the integers is straightforward, in that it merely involves tossing in $\{-1, -2, -3, \ldots\}$, where these additive inverses are simply declared to exhibit the behavior of additive inverses. However, the extension of the integers to the rationals is a different sort of program. The set of rational numbers is not a superset of the integers in the same way that the integers represent a superset of the whole numbers. Rational numbers are conceived as integer pairs p/q, where q is nonzero, and rational equality, addition, and multiplication are created from scratch, using equality, addition, and multiplication of integers as ingredients.

By creating the rational numbers from the integers as we did, we say that the integers can be *isomorphically embedded* in the rationals. Given an integer p, we may replace p with the rational expression $p/1$. We then can see that rational addition of these replacement forms behaves in a way that corresponds to integer addition:

$$p/1 \oplus q/1 \equiv (p \cdot 1 + 1 \cdot q)/(1 \cdot 1) \equiv (p + q)/1$$

where $(p + q)/1$ corresponds to $p + q$ in the integers. A similar result holds for multiplication:

$$(p/1) \otimes (q/1) \equiv (pq)/(1 \cdot 1) \equiv (pq)/1$$

The rational numbers therefore comprise a smallest possible extension of $\{0, 1\}$ that has closure of addition and multiplication, existence of additive and multiplicative identities, and closure under additive and multiplicative inverses. In other words, the rationals are the most efficient possible extension of $\{0, 1\}$ that satisfies assumptions A1–A19.

Up to this point, we have said virtually nothing about *irrational numbers*—real numbers that are not rational. You have worked with lots of numbers that are irrational: π, e, and most numbers of the form $\sqrt[n]{x}$. You have also been taught that irrational numbers have a decimal representation that does not terminate and does not fall into a pattern of repetition. However, you might never have seen a definitive argument that any particular number is irrational. In this section we address the existence of irrational number. But first, a little history.

Some of the intuitive properties of the real numbers that we likely take for granted were not a part of the thinking of the ancient Greeks, most notably the Pythagoreans, that very secret order of thinkers who produced some amazingly sophisticated mathematics. One such intuitive feel we probably have is that the set of real numbers is a sort of continuum, a set that can be visualized in terms of the ordered points on a straight line, where rationals and irrationals are spread up and down like strewn grains of salt and pepper, leaving no holes in the continuum. The Greeks had no concept of irrational numbers at first, and now is a good time to touch briefly on some of their views.

To the Greeks, numbers were conceived in terms of the geometric concepts of length, area, and volume, all of which were *constructible* according to certain idealized rules. They had very sophisticated techniques for imagining their idealized construction using the only two geometric shapes they considered perfect: straight line segments (crudely constructible with a straight-edge) and circles (crudely constructible with a compass). In his all-important compilation of the best of Greek mathematics (a five-volume set called *Elements*), Euclid demonstrated amazingly sophisticated mathematical results, such as the theorem attributed to Pythagoras, using only some assumptions about circles and line segments.

With these techniques of construction, it is possible to imagine the drawing, if you will, of a line segment whose length is any positive rational number in the following way. Beginning with a line segment whose length is arbitrarily declared to be one unit, it is possible to construct segments representing 2, 3, 4, and so on, by extending the given segment with a straightedge, then twirling a compass around to tack the measured unit length onto itself end to end. It is also pretty easy to take a segment of length n, and use similar triangles to construct a line segment of length $1/n$. (See Figure 3.11.) With another technique by which segments of length a and b can be used to construct a segment of length ab, all the positive rational lengths are constructible.

Exercise 3.10.1 Devise a technique similar to that described in Figure 3.11 by which a segment of length ab can be constructed from segments of length a and b.

The fact that segments of arbitrarily long integer length n could be constructed meant that segments of arbitrarily short length $1/n$ could also be constructed by the reciprocation technique. Such a short line segment could then, in theory, be added to itself as many times as needed to produce a segment of any arbitrary length m/n. The Greeks erroneously assumed the converse—that a segment of any constructible length could be constructed by imagining it to be a finite sum of very short segments of the form $1/n$. Slap any segment

1. Construct segment *AB* of length *n*.
2. Construct segment *AC* of length 1.
3. Construct segment *BC*.
4. Construct segment *FE* parallel to *BC*.
5. Segment *AE* has length 1/*n*.

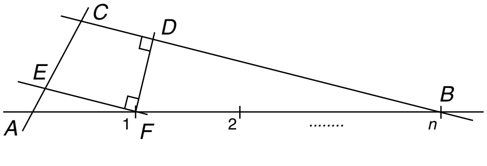

Figure 3.11 Constructing length $1/n$ using similar triangles.

down onto the imaginary paper using techniques of construction. How long is it? Some whole multiple of a (possibly very short) segment that can be constructed by reciprocating a segment of positive integer length n. In modern language, what number is represented by the length of a segment? It is always a rational one.

Now suppose we are given two segments of arbitrary lengths. If both of them can be visualized in this way, what sort of relationship must exist between these two segments? In the same way that we would find a common denominator of two fractions m/n and p/q, they would say that there is a single segment, perhaps very short, that can be attached to itself a finite number of times to produce each of the two given segments. In our language, $1/nq$ can be added to itself mq times to produce m/n and np times to produce p/q. This assumed relationship between two segments is called *commensurability*. The Greeks believed that all constructible segments are commensurable, which is equivalent to our believing that all real numbers are rational.

There was a problem with this, and the Greeks eventually figured it out. Take a segment of length one, then construct another segment of length one at one endpoint and at a right angle to this first segment. Then sketch the hypotenuse. This hypotenuse line segment is obviously constructible. We just described how to do it easily. Furthermore, by the Pythagorean theorem, its length d satisfies $d^2 = 2$. The problem is that it is not commensurable with some other constructible segments. That is, d is not rational. This is the amazing result that the Pythagoreans discovered, and it created no small crisis. Given that the theorem traditionally named after Pythagoras was already known, it boiled down to this fact: Since "the sum of the [areas of the] squares on the legs of a right triangle is equal to the [area of the] square on the hypotenuse" (see Fig. 3.12), then, as we would write it, $\sqrt{2}$ is a constructible length. Therefore, $\sqrt{2}$ is commensurable with 1, or, in our language, $\sqrt{2}$ must be writable in the form p/q, where p and q are integers and q is nonzero. The traditional proof that this is impossible is attributed to Aristotle. And, by the way, it is the one you will discover in the following exercise, with a few hints.

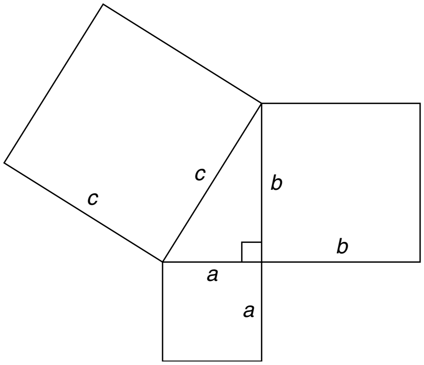

Figure 3.12 A sketch of the Pythagorean theorem.

EXERCISE 3.10.2 $\sqrt{2}$ is irrational.[33,34,35,36,37]

A few words are in order about irrational numbers. Let's work backwards from Exercise 3.10.2. Exercise 3.10.2 says $\sqrt{2}$ is not rational, while Theorem 3.9.1 says $\sqrt{2}$ is real. Thus there do, in fact, exist real numbers that are irrational. But Theorem 3.9.1 is based on assumption A22, which we have said very little about. As we said, assumption A22 is not an axiom of the real numbers. It can be proved from the Least Upper Bound axiom, which is a standard axiom of the real numbers. So what we discover is that the irrational numbers owe their existence to the Least Upper Bound axiom.

EXERCISE 3.10.3 The proof of Exercise 3.10.2 can be easily generalized to demonstrate that $\sqrt{p}$ is irrational for any prime number p. Explain how your proof of Exercise 3.10.2 can be adapted into a proof of this more general claim.[38]

EXERCISE 3.10.4 The sum of a rational and an irrational is irrational.[39]

EXERCISE 3.10.5 The product of a nonzero rational and an irrational is irrational.

EXERCISE 3.10.6 Let $a, b, c,$ and d be rational numbers, and suppose $a + b\sqrt{2} = c + d\sqrt{2}$. Does it follow that $a = c$ and $b = d$?[40]

3.11 Relations in General

Equivalence relations are just one example of the more general mathematical idea of a *relation*. A relation is a set construction that puts all kinds of element comparisons such as equality, less than, divisibility, and subset inclusion into one mathematical idea. It is also a way of linking elements of two different sets together, and it is a context in which functions can be defined. In this section, we define a relation and look at examples and special kinds of relations. But first, we define the *Cartesian product* of two sets A and B by

$$A \times B = \{(a, b) : a \in A, b \in B\} \tag{3.64}$$

[33] Naturally, you will want to assume $\sqrt{2}$ is rational and arrive at a contradiction.

[34] Write $\sqrt{2} = m/n$, where you can assume no common factors between m and n.

[35] Square both sides and look at the contrapositive of Corollary 2.4.5 .

[36] If m^2 is even then, ... What does this mean? Cancel out a 2.

[37] So n^2 is even. What contradiction does this cause?

[38] Corollary 2.5.12 might help.

[39] See Exercise 1.2.18(i).

[40] If $b \neq d$, then what must be true of $\sqrt{2}$?

the set of ordered pairs where the first term in the pair is an element of A and the second is an element of B. A *relation from A to B* is defined to be any subset of $A \times B$.

Example 3.11.1 For $A = \{1, 2, 3\}$ and $B = \{11, 12, 13, 21, 22, 23, 31, 32, 33\}$, the set

$$R = \{(1, 11), (2, 21), (2, 22), (3, 31), (3, 32), (3, 33)\} \tag{3.65}$$

is a relation from A to B. ■

In this section, we are going to delve only into subsets of $A \times A$, which we call a *relation on A* instead of a relation from A to A.

Example 3.11.2 Let $A = \{1, 2, 3, 4, 5, 6\}$. Then

$$R = \{(1, 3), (1, 5), (2, 4), (2, 6), (3, 5), (4, 6)\} \tag{3.66}$$

is a relation on A. ■

You might have seen the Cartesian plane written as $\mathbb{R} \times \mathbb{R}$, or more succinctly as $\mathbb{R}^2$. If R is a relation on the real numbers, we can represent it graphically as a set of points in the xy-plane.

Example 3.11.3 $\{(x, y) : y > |x| + 1\}$ is a relation on the real numbers (Fig. 3.13). ■

EXERCISE 3.11.4 For each of the following relations on the real numbers, sketch the set of included points in the xy-plane.

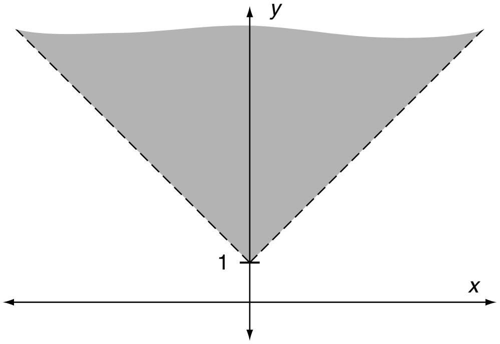

Figure 3.13 The relation $y > |x| + 1$.

(a) $\{(x, y) : x \leq y\}$

(b) $\{(x, y) : x^2 + y \leq 4\}$

(c) $\{(x, y) : |x| < 1\}$

(d) $\{(x, y) : x^2 + y^2 \geq 1\}$

You might think there is nothing particularly exciting about a relation, since any collection of ordered element pairs qualifies as a relation. What makes the idea take shape and become mathematically important is that we can lay out certain criteria by which pairs are included in the relation, and these different criteria might have particular properties that make interesting statements about A.

Example 3.11.5 Define a relation $R \subseteq \mathbb{Q} \times \mathbb{Q}$ by $R = \{(p/q, r/s) : ps = qr\}$. This relation consists of all pairs of rational numbers that are equivalent, as was defined in Definition 3.8.2. So $(3/8, -30/-80)$ and $(0/2, 0/12)$ are in R, while $(2/5, 5/3)$ is not. ■

Example 3.11.6 $R_1 = \{(a, b) : a < b\}$ and $R_2 = \{(a, b) : a \mid b\}$ are relations on the integers. ■

Example 3.11.7 For a set $A, \emptyset$ and $A \times A$ are relations on A. ■

The *power set* of a set A is defined to be the family of all subsets of A and is written $\mathcal{P}(A)$. For example, if $A = \{1, 2, 3\}$, the power set of A is

$$\mathcal{P}(A) = \{\emptyset, \{1\}, \{2\}, \{3\}, \{1, 2\}, \{1, 3\}, \{2, 3\}, \{1, 2, 3\}\} \tag{3.67}$$

EXERCISE 3.11.8 List all elements of the power set of the following sets.

(a) $\emptyset$

(b) $\{1\}$

(c) $\{1, 2\}$

(d) $\{1, 2, 3, 4\}$

Example 3.11.9 For $A = \{1, 2, 3\}$, define a relation on the power set of A by $R = \{(A_1, A_2) : A_1 \subseteq A_2\}$. Thus $(\{1\}, \{1\})$ and $(\emptyset, A)$ are in R, but $(A, \{2\})$ and $(\{1, 2\}, \emptyset)$ are not.

In the previous examples, a particular pair (x, y) is included in the relation if some stated property $P(x, y)$ is true. That is, $(x, y) \in R$ if x and y are *related* according to some criterion. We give names to some relations when the criterion for inclusion of ordered pairs in R has certain properties. ■

Definition 3.11.10 Suppose $R \subseteq A \times A$ has the following properties.

(E1) For all $x \in A$, $(x, x) \in R$. (Reflexive)

(E2) For all $x, y \in A$, if $(x, y) \in R$, then $(y, x) \in R$. (Symmetric)

(E3) For all $x, y, z \in A$, if $(x, y), (y, z) \in R$, then $(x, z) \in R$. (Transitive)

Then R is called an *equivalence relation* on A.

Notice how E1–E3 here say the same things as E1–E3 from Definition 3.6.1, but in different language. Instead of saying that $x \equiv x$ for all $x \in A$ is a true statement, we say that all ordered pairs of the form (x, x) are in the relation. Instead of saying that the truth of $x \equiv y$ implies the truth of $y \equiv x$, now we say that the inclusion of the ordered pair (x, y) in the relation is always accompanied by the inclusion of (y, x). And instead of saying $x \equiv y$ and $y \equiv z$ implies $x \equiv z$, we say that the presence of (x, y) and (y, z) in the relation is always accompanied by the presence of (x, z). Thus an equivalence relation on A is a subset of $A \times A$ with some special properties that motivate a partition of A.

Anytime we define a term that describes how elements of a set may or may not compare to each other, we can use that basis of comparison as a criterion for inclusion in a relation. For example, the symbol $\leq$ denotes a way that two real numbers relate or compare to each other. In Exercise 3.11.4(a), the statement $x \leq y$ is a criterion for the inclusion of a point (x, y) in a relation. The relation defined by $\leq$ has some special properties that motivate a definition of another type of relation. Since it is common to write $x\mathrm{R}y$ as a quicker way to write $(x, y) \in R$, and to say that x is *related to* y instead of saying (x, y) is an element of the relation, we use this somewhat more efficient notation in the next definition and example.

Definition 3.11.11 Suppose R is a relation on a set A with the following properties.

(O1) For all $x \in A$, $x\mathrm{R}x$. (Reflexive)

(O2) For all $x, y \in A$, if $x\mathrm{R}y$ and $y\mathrm{R}x$, then $x = y$. (Antisymmetric)

(O3) For all $x, y, z \in A$, if $x\mathrm{R}y$ and $y\mathrm{R}z$, then $x\mathrm{R}z$. (Transitive)

Then R is called an *order relation* on A, or a *partial ordering* of A.

Example 3.11.12 Show that the relation in Exercise 3.11.4(a) is an order relation on the real numbers.

Solution We show that R has properties O1–O3.

(O1) Since $x = x$ for all $x \in \mathbb{R}$, we have that $x \leq x$, so that $x\mathrm{R}x$ for all x.

(O2) Suppose $x\mathrm{R}y$ and $y\mathrm{R}x$. Then $x \leq y$ and $y \leq x$, so that $x = y$.

(O3) Suppose $x\mathrm{R}y$ and $y\mathrm{R}z$. Then $x \leq y$ and $y \leq z$. Now if $x = y$, then $z - x = z - y \geq 0$, so that $x \leq z$. Similarly, if $y = z$, we have that $x \leq z$. If $x < y$ and $y < z$, then $x < z$ by Exercise 2.2.5(e). In any case, $x\mathrm{R}z$.

Since R has properties O1–O3, it defines an order relation on the real numbers. ■

We can be even more efficient with our language and notation than we were in Example 3.11.12 by dispensing with the statement $x\mathrm{R}y$ and saying simply that $\leq$ defines an order relation on the real numbers. That would make the verification above look like this.

Solution We show that $\leq$ has properties O1–O3.

(O1) For all $x \in \mathbb{R}$, $x = x$ so that $x \leq x$.

(O2) If $x \leq y$ and $y \leq x$, then $x = y$.

(O3) Suppose $x \leq y$ and $y \leq z$. If $x = y$, then $z - x = z - y \geq 0$, so that $x \leq z$. Similarly, if $y = z$, we have that $x \leq z$. If $x < y$ and $y < z$, then $x < z$ by Exercise 2.2.5(e). In any case, $x \leq z$.

Since $\leq$ has properties O1–O3, it defines an order relation on the real numbers. ■

Verifying the claim in the next exercise will be quick if you will reference the applicable results from Section 3.2.

Exercise 3.11.13 If A is any set, then $\subseteq$ defines an order relation on $\mathcal{P}(A)$.

For the set $\{1, 2, 3\}$, the order relation defined by $\subseteq$ can be illustrated with a *directed graph* as in Figure 3.14. The arrow from $\{1\}$ to $\{1, 2\}$ indicates that the former is related to the latter by $\subseteq$. Notice that $\{1, 2, 3\}$ is reachable from $\{1\}$ by

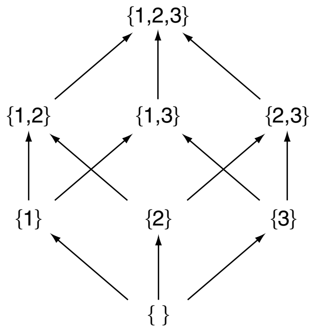

Figure 3.14 Directed graph of the partial order relation $\subseteq$ on $\mathcal{P}(\{1, 2, 3\})$.

a *directed path* by way of either $\{1, 2\}$ or $\{1, 3\}$, and either of these directed paths indicates that $\{1\} \subseteq \{1, 2, 3\}$.

EXERCISE 3.11.14 Divisibility defines an order relation on the positive integers.

EXERCISE 3.11.15 Sketch a directed graph to illustrate the order relation of divisibility on $\{1, 2, 3, 4, 6, 8, 9, 12, 18, 24, 27, 36, 54\}$.

Another type of relation is reminiscent of $<$. We use the expression $x \not\mathrm{R} y$ to mean $(x, y) \notin R$.

Definition 3.11.16 Suppose R is a relation on a set A with the following properties.

(S1) For all $x \in A$, $x \not\mathrm{R} x$. (Irreflexive)

(S2) For all $x, y \in A$, if $x\mathrm{R}y$, then $y \not\mathrm{R} x$. (Asymmetric)

(S3) For all $x, y, z \in A$, if $x\mathrm{R}y$ and $y\mathrm{R}z$, then $x\mathrm{R}z$. (Transitive)

Then R is called a *strict order relation* on A, or a *strict ordering* of A.

Example 3.11.17 Show that $<$ defines a strict order relation on the integers.

Solution We show that $<$ satisfies properties S1–S3.

(S1) Since $x - x = 0$, we have that $x - x \not< 0$, so $x \not< x$ for all integers x.

(S2) Suppose $x < y$. Then $y - x > 0$, so that $y - x \not< 0$. Thus $y \not< x$.

(S3) By Exercise 2.2.5(e), if $x < y$ and $y < z$, then $x < z$. ■

EXERCISE 3.11.18 We may use the results of Example 3.11.17 to define a relation on the rational numbers. In this exercise, if p/q and r/s are rational numbers, we assume that q and s are both positive. Define $<_{\mathbb{Q}}$ to be a relation on the rational numbers, where $p/q <_{\mathbb{Q}} r/s$ provided the integer inequality $ps < qr$ is satisfied. Use the fact that $<$ is a strict order relation on the integers to show that $<_{\mathbb{Q}}$ is a strict order relation on the rationals.

EXERCISE 3.11.19 For any set A, $\subset$ defines a strict order relation on $\mathcal{P}(A)$.

One notable difference between $\leq$ on the integers and divisibility on the positive integers is that any pair of integers we choose is comparable in one way or the other by $\leq$. That is, for all integers a and b, either $a \leq b$ or $b \leq a$. However, for divisibility, there are many pairs of positive integers that are not comparable at all. For example, neither $6 \mid 10$ nor $10 \mid 6$ is true. This distinction motivates yet another type of relation.

Definition 3.11.20 Suppose R is a relation on a set A with the following properties.

(T1) For all $x \in A$, $x\text{R}x$. (Reflexive)

(T2) For all $x, y \in A$, if $x\text{R}y$ and $y\text{R}x$, then $x = y$. (Antisymmetric)

(T3) For all $x, y, z \in A$, if $x\text{R}y$ and $y\text{R}z$, then $x\text{R}z$. (Transitive)

(T4) For all $x, y \in A$, either $x\text{R}y$ or $y\text{R}x$. (Totality)

Then R is called a *total order relation* on A, or a *total ordering* of A.

Notice that properties T1–T3 are the same as O1–O3, so a total order relation is a special kind of order relation where every pair of elements is comparable in one way or another.

Example 3.11.21 The relation $\leq$ is a total ordering of the real numbers. ■

Example 3.11.22 For $A = \{1, 2, 3\}$, $\subseteq$ does not define a total order of $\mathcal{P}(A)$, for $\{1, 2\} \not\subseteq \{2, 3\}$ and $\{2, 3\} \not\subseteq \{1, 2\}$. ■

One of our assumptions from Chapter 0 is the well-ordering principle. To say that a is the smallest element of a non-empty subset of the whole numbers is to say that $a \leq x$ for all x in the set. The last type of relation we define is a generalization of the well-ordering principle.

Definition 3.11.23 Suppose R is a relation on a set A with the following properties.

(W1) For all $x \in A$, $x\text{R}x$. (Reflexive)

(W2) For all $x, y \in A$, if $x\text{R}y$ and $y\text{R}x$, then $x = y$. (Antisymmetric)

(W3) For all $x, y, z \in A$, if $x\text{R}y$ and $y\text{R}z$, then $x\text{R}z$. (Transitive)

(W4) If S is any non-empty subset of A, then there exists $a \in S$ such that $a\text{R}x$ for all $x \in S$. (Well ordering)

Then R is called a *well order relation* on A, or a *well ordering* of A.

Because properties W1–W3 are the same as O1–O3, a well ordering of A is an order relation. The feature that property W4 adds is that every non-empty subset of A has what we might call a least element.

Exercise 3.11.24 The least element of a well-ordered set is unique.

Exercise 3.11.25 Show that a well ordering is a total ordering by showing that property W4 implies T4.

Though a well ordering is a total ordering, a total ordering is not necessarily a well ordering. For example, $\leq$ is a total ordering of the integers that is not a well ordering, for the entire set of integers does not contain a least element. Since W4 implies T4, but not vice versa, property W4 is therefore stronger than T4.

Just because a given a set with an order relation fails to be a well ordering, it does not mean that the set cannot be well ordered by some other relation. In fact, one of the most notable results in the modern theory of sets is that *any* set can be well ordered. Ernst Zermelo demonstrated this in 1904, using an axiom of set theory that we have said nothing about so far in this text. The *axiom of choice* is a somewhat mysterious axiom of set theory that is simple to state but seldom addressed at this level of the mathematical game. It says, "Given any family $\mathcal{F}$ of mutually disjoint non-empty sets, there is a set S that contains a single element from each set in $\mathcal{F}$." Thus S can be thought of as the result of having chosen a distinct representative element from each set in $\mathcal{F}$. In the axiom of choice, there is enough strength to demonstrate that for any set, there exists a relation on $A \times A$ that is a well ordering of A. One way to see how the integers can be well ordered is to list its elements as $\langle 0, -1, 1, -2, 2, -3, 3, -4, 4, \ldots \rangle$ and define $x\mathrm{R}y$ if x does not come after y in this listing. When the integers are ordered in this way, every non-empty subset has a least element.

Functions

Second only to sets, *functions* are likely the most important mathematical concept. Interestingly, functions can be defined solely in terms of sets, and some authors take that very formal approach. In this chapter, we define functions in a somewhat more informal way, and then we study some of the most important basic results about functions.

4.1 Definition and Examples

In Section 3.11 we defined a relation from A to B as any set of ordered pairs (x, y), where $x \in A$ and $y \in B$. A *mapping* from A to B is another term for relation, though it conjures up the imagery that, given a particular element of A, some process or computation is done on it to pair it with an element or elements of B. We use this input–output imagery to define a *function* as a mapping with two special features.

Definition 4.1.1 Given two non-empty sets A and B, a *function* f is a rule, or set of instructions, by which each element of A is paired with exactly one element of B. Set A is called the *domain* and is denoted $\operatorname{Dom} f$. Set B is called the *codomain*. Notationally, if f is a function from A to B, we write $f : A \to B$. If $x \in A$ is paired by f with $y \in B$, we write $f(x) = y$ or $x \mapsto y$. We say that y is the *image* of x, or that x *maps* to y, and x is a *pre-image* of y. The subset of B consisting of the images of all elements of A is called the *range* and is denoted $\operatorname{Rng} f$. (See Figure 4.1 for a basic sketch.)

Let's sort out the details of Definition 4.1.1. First, for a mapping $f : A \to B$ to be a function, *every* $x \in A$ must have some image $y \in B$. That is

(F1) For every $x \in A$, there exists $y \in B$ such that $f(x) = y$.

In addition to property F1, the rule defining f must produce a unique $f(x)$ for all $x \in A$. We say that f is *well defined* if $f(x)$ is unique for every $x \in A$. How can

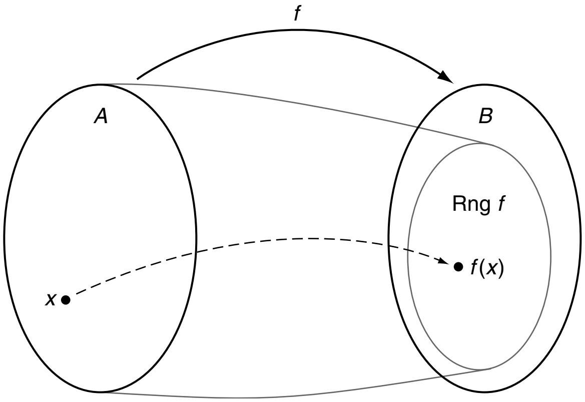

Figure 4.1 Schematic of a function: $f : A \to B$.

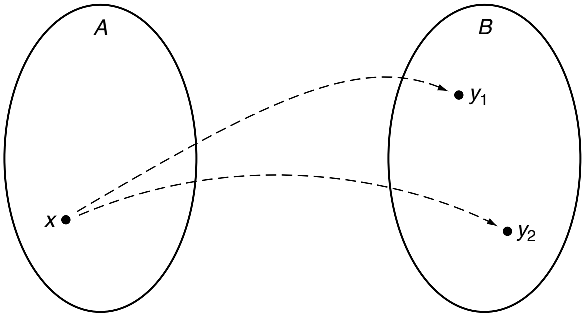

Figure 4.2 The image of x is not unique.

we say mathematically that the image of $x \in A$ must be unique? (See Figure 4.2 for a sketch of what must not happen.) One way to say that the rule for f cannot produce multiple y-values for a single x is to say

(F2) If $y_1, y_2 \in B$ are such that $f(x) = y_1$ and $f(x) = y_2$, then $y_1 = y_2$.

At first glance, F2 might seem like a statement that would be true for any mapping. After all, if $f(x) = y_1$ and $f(x) = y_2$, then the transitive property of equality ought to allow us to say $y_1 = y_2$. Unfortunately, this way of stating F2 does not reveal potential problems in the way f is defined. Here are two examples that illustrate how a mapping can fail to be well defined.

Example 4.1.2 It is possible that the rule defining $f : A \to B$ can be ambiguous. For example, let $f : \mathbb{R}^+ \to \mathbb{R}$ be defined in the following way. For $x > 0$, define $f(x)$ to be a solution y to the equation $y^2 = x$. You cannot doubt that this set of instructions generates a real number $f(x)$ from any x in the domain, so that property F1 is satisfied. However, the rule is ambiguous for every x in the domain. Letting $x = 25$, $y_1 = 5$ and $y_2 = -5$, we have demonstrated the existence of y_1 and y_2 in the codomain where $f(x) = y_1$ and $f(x) = y_2$, but where $y_1 \neq y_2$. ■

Example 4.1.2 illustrates a possible pitfall you might run into with the notation $f(x)$ if you have not verified that f is well defined. Unless you know that f is well defined, you might find yourself using an expression like $f(25)$ as one name for more than one distinct thing. We must make sure that the set of instructions for generating $f(x)$ does not produce more than one value for any x in the domain.

Here is another example of how a mapping can fail to be well defined, this time because a single element in the domain might have more than one distinct name.

Example 4.1.3 Define $f : \mathbb{Q} \to \mathbb{Z}$ in the following way. For a rational number p/q, define $f(p/q) = p$. That is, the image of a rational number written in a standard form of integer over integer is defined to be the numerator. Unlike Example 4.1.2, there is no ambiguity concerning what $f(p/q)$ is. The problem here is that the standard definition of equality in the rationals (Section 3.8.1) lumps a lot of different expressions of the form p/q into the same equivalence class. As a specific example, though $2/5 = 6/15$, $f(2/5) \neq f(6/15)$. The fact that one domain element can be addressed by more than one name causes problems, because the image of an element of the domain depends on the form in which it is written. Thus f is not well defined because we have exhibited an x in the domain and y_1 and y_2 in the codomain where $f(x) = y_1$, $f(x) = y_2$, but $y_1 \neq y_2$. ■

Example 4.1.3 illustrates a good way to restate property F2. If x_1 and x_2 are two names for the same thing, that is, if $x_1 = x_2$, then the rule must produce the same functional value for x_1 and x_2. That is, $f(x_1) = f(x_2)$.

(F2) If $x_1, x_2 \in A$ and $x_1 = x_2$, then $f(x_1) = f(x_2)$.

Suppose someone gives us two sets A and B, and a rule f for pairing elements of A with elements of B. If we are asked to verify that f is a function, then we must verify that F1–F2 hold.

Example 4.1.4 Show that $f : \mathbb{R} - \{1\} \to \mathbb{R}$ defined by

$$f(x) = \frac{x^2 + 2}{x - 1}$$

is a function.

Solution

(F1) Pick any $x \in \mathbb{R} - \{1\}$. Then by properties A3 and A9 (closure of addition and multiplication, respectively), $x^2 + 2$ and $x - 1$ are both real numbers. Furthermore, since $x \neq 1$, then $x - 1$ is nonzero, so that $(x - 1)^{-1}$ exists as a real number by property A13. Finally, $(x^2+2)\cdot(x-1)^{-1} \in \mathbb{R}$ by closure of multiplication. Thus $f(x)$ exists in $\mathbb{R}$.

(F2) Pick $x_1, x_2 \in \mathbb{R}-\{1\}$ and suppose $x_1 = x_2$. Since addition and multiplication are well defined, it follows that $x_1^2+2 = x_2^2+2$ and $x_1-1 = x_2-1$. By the uniqueness of multiplicative inverses, $(x_1-1)^{-1} = (x_2-1)^{-1}$. Again, since multiplication is well defined, $(x_1^2+2)/(x_1-1) = (x_2^2+2)/(x_2-1)$, so that $f(x_1) = f(x_2)$. ■

Example 4.1.4 illustrates a broad result that is pretty easy to see through the example alone. If f is a rule that can be written in the form of polynomial over polynomial, that is, if

$$f(x) = \frac{a_n x^n + a_{n-1}x^{n-1} + \cdots + a_1 x + a_0}{b_m x^m + b_{m-1}x^{m-1} + \cdots + b_1 x + b_0} \tag{4.1}$$

where all the a_i and b_j are real numbers, then assuming we avoid values of x that make the denominator zero, f is a function from an appropriate subset of the real numbers into the real numbers. Here is another way to say it. Since addition and multiplication are well defined in the real numbers, and since additive and multiplicative inverses are unique, then any mapping whose rule is built up from the operations of $+, -, \times, \div$ and whose domain avoids division by zero will be well defined.

We can go even further. Thanks to Theorem 3.9.1 and Definition 3.9.3, we have ensured that $\sqrt[n]{x}$ is well defined for all real x if n is odd, and for all nonnegative x if n is even. Therefore, we can say that any mapping whose rule is built up algebraically with $+, -, \times, \div, \square^n$, and $\sqrt[n]{\ }$, and whose domain is a subset of the real numbers that avoids division by zero and even roots of negative numbers will be well defined. For example,

$$f(x) = 5 + x + \sqrt[3]{\frac{1+\sqrt{x^3 - 1/(x-3)}}{x^5 - 1 + x^{-5}}} \tag{4.2}$$

is a function whose domain is some hideous subset of the real numbers. Functions built up as is f in Eq. (4.2) are called *algebraic*.

There is a simple function that will come in handy in Section 4.6. In its simplicity, it illustrates pretty well the sorts of things we have to show when verifying a mapping is a function and when showing a function has certain properties.

Example 4.1.5 Let m and n be integers, where n is positive, and write $\mathbb{N}_n = \{1, 2, \ldots, n\}$ and $\mathbb{N}_n^m = \{1+m, 2+m, \ldots, n+m\}$. Define a translation mapping $T : \mathbb{N}_n \to \mathbb{N}_n^m$ by $T(x) = x + m$. Show T is a function.

Solution We show that T satisfies properties F1–F2.

(F1) Pick $x \in \mathbb{N}_n$. Since $1 \leq x \leq n$, it follows that $1 + m \leq x + m \leq n + m$, or $1 + m \leq T(x) \leq n + m$. Thus $T(x) \in \mathbb{N}_n^m$.

(F2) Pick $x_1, x_2 \in \mathbb{N}_n$ and suppose $x_1 = x_2$. Then since addition is well defined, $T(x_1) = x_1 + m = x_2 + m = T(x_2)$. Thus T is well defined. ■

Notice that showing F1 in Example 4.1.5 involves not only showing that $T(x)$ exists, but that it lies in the codomain.

If a and b are real numbers such that $a < b$, then we use the following notation to represent certain subsets of the real numbers:

$$\begin{aligned}
[a, +\infty) &= \{x \in \mathbb{R} : x \geq a\} \\
(a, +\infty) &= \{x \in \mathbb{R} : x > a\} \\
(-\infty, b] &= \{x \in \mathbb{R} : x \leq b\} \\
(-\infty, b) &= \{x \in \mathbb{R} : x < b\} \\
[a, b] &= \{x \in \mathbb{R} : a \leq x \leq b\} \\
(a, b) &= \{x \in \mathbb{R} : a < x < b\}
\end{aligned}$$

EXERCISE 4.1.6 Show that each of the following are functions by showing they have properties F1–F2.

(a) $f : [0, 1] \to [0, 10]$ defined by $f(x) = 3x + 4$

(b) $f : \mathbb{R} \to [-2, \infty)$ defined by $f(x) = x^2 - 1$

(c) $f : \mathbb{R} \to \mathbb{R}$ defined by $f(x) = \sqrt[3]{1 - x}$[1]

If $f(x)$ is a formula used to define a function whose domain and codomain are some subsets of the real numbers, then the *natural domain* of f is defined to be the largest subset of the real numbers for which $f(x)$ exists as a real number. For example, if $f(x) = x/(x + 2)$, then the rule for f can be used to generate a real number for all $x \neq -2$. Thus the natural domain of f is all real numbers except -2.

EXERCISE 4.1.7 Determine the natural domain of each of the following functions.

(a) $f(x) = \dfrac{3x + 5}{2x - 1}$

(b) $f(x) = \sqrt{x + 4}$

[1] See the comments after Definition 3.9.3.

(c) $f(x) = \dfrac{1}{\sqrt{x+4}}$

(d) $f(x) = \sqrt{4 - x^2}$

(e) $f(x) = \dfrac{x-1}{x^3 - x}$

EXERCISE 4.1.8 Let A be any non-empty set, and define the *identity* mapping $i : A \to A$ by $i(x) = x$ for all $x \in A$. Show that the identity mapping is a function.

Now we want to define a notion of equality for two functions. Certainly, we want the definition to be an equivalence relation. The standard definition of function equality is the following.

Definition 4.1.9 Two functions f and g are said to be equal provided they have the same domain and codomain, and $f(x) = g(x)$ for all x in the domain.

To show that Definition 4.1.9 is an equivalence relation is not difficult. For example, to show E1, we note that f has the same domain and codomain as itself, and $f(x) = f(x)$ for all x in the domain. Thus $f = f$.

EXERCISE 4.1.10 Finish the proof that Definition 4.1.9 is an equivalence relation by showing it has properties E2–E3.

At the beginning of this section, we said that mappings and relations are synonymous, even though they evoke different sorts of imagery. Some authors define a function as a special kind of relation, that is, a subset of $A \times B$ with certain properties. Instead of imagining elements of A being associated with elements of B via the input–output imagery of Definition 4.1.1, it is possible to define a function as a subset of $A \times B$ with two additional restrictions that coincide with properties F1–F2.

For a relation $R \subseteq A \times B$ to be a function, property F1 says that, for all $x \in A$, there must exist $y \in B$ such that $(x, y) \in R$. Property F2 says that if $(x, y_1), (x, y_2) \in R$, then $y_1 = y_2$.

EXERCISE 4.1.11 Let $A = \{1, 5, 10\}$ and $B = \{1, 2, 8, 9\}$. For each of the following relations, state whether it is a function.

(a) $\{(1, 1), (5, 9)\}$

(b) $\{(1, 2), (10, 1), (5, 9)\}$

(c) $\{(5, 1), (1, 1), (10, 2), (1, 8)\}$

(d) $\{(5, 2), (10, 2), (1, 2)\}$

4.2 One-to-one and Onto Functions

Property F1 says that a function $f : A \to B$ must map every element of A to some element of B. Property F2 says that this image must be unique, so that the path from x to $f(x)$ must not have a fork in the road (Figure 4.2). Some functions have features that are analogous to F1–F2 when viewed, as it were, in the reverse direction. Analogous to every element of A having an image in B, perhaps every element of B has a pre-image in A. A function with this special feature is said to be *onto*, so that Rng $f = B$. Also, analogous to the uniqueness of the image of every element of A, perhaps every element of Rng f has a unique pre-image. This type of function is called *one-to-one*. In a one-to-one function, the forbidden merging of the path from different elements of A to the same image in B is illustrated in Figure 4.3. In this section, we define these terms and look at features of functions that are one-to-one and/or onto.

Definition 4.2.1 Suppose $f : A \to B$ is a function. Then f is said to be *onto* provided for every $b \in B$, there exists $a \in A$ such that $f(a) = b$. That is,

$$f \text{ is onto } \Leftrightarrow (\forall y \in B)(\exists x \in A)(y = f(x)) \tag{4.3}$$

As a notational shorthand, we write $f : A \xrightarrow{\text{onto}} B$. An onto function is also called a *surjection*.

Contrast the definition of onto in (4.3) with property F1, which says

$$(\forall x \in A)(\exists y \in B)(y = f(x))$$

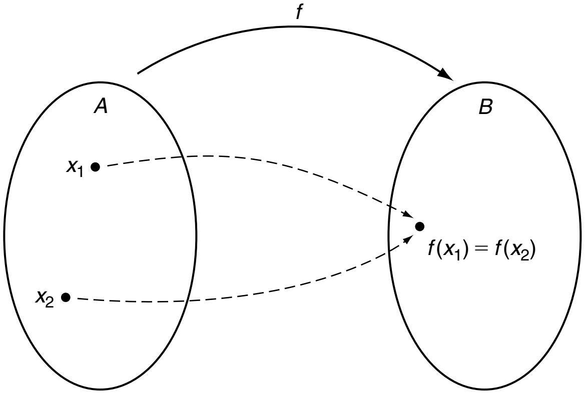

Figure 4.3 A function that is not one-to-one.

Example 4.2.2 Show that $f : \mathbb{R} \to \mathbb{R}$ defined by $f(x) = x^3 - 1$ is a surjection.

Solution You may recall from pre-calculus that the graph of a polynomial function of odd degree extends all the way up and down the y-axis, so that any value chosen on the y-axis has a pre-image on the x-axis. But how do we prove f is onto? We have to choose an arbitrary y in the codomain and work backwards to find some x in the domain that maps to it. In other words, we must find a value of x that makes the equation $y = f(x)$ true. Here is the proof.

Pick any real number y, and let $x = \sqrt[3]{y+1}$. By Theorem 3.9.1, x is a real number, and

$$f(x) = f(\sqrt[3]{y+1}) = [\sqrt[3]{y+1}]^3 - 1 = y + 1 - 1 = y \tag{4.4}$$

Thus f is onto. ■

EXERCISE 4.2.3 Show that the function $T : \mathbb{N}_n \to \mathbb{N}_n^m$ from Example 4.1.5 is a surjection.

EXERCISE 4.2.4 Let $f : [0, 2] \to [1, 7]$ be defined by $f(x) = 3x + 1$. Show that f is a surjection.

EXERCISE 4.2.5 Let A be a non-empty set, and $i : A \to A$ the identity function (Exercise 4.1.8). Show that i is onto.

EXERCISE 4.2.6 What does it mean to say that a function is not onto?

EXERCISE 4.2.7 Show that $f : [-2, 3] \to [0, 10]$ defined by $f(x) = x^2 + 1$ is not onto.

Definition 4.2.8 Suppose $f : A \to B$ is a function. Then f is said to be one-to-one provided $f(x_1) = f(x_2)$ implies $x_1 = x_2$ for all $x_1, x_2 \in A$. That is,

$$f \text{ is one-to-one } \Leftrightarrow (\forall x_1, x_2 \in A)([f(x_1) = f(x_2)] \to [x_1 = x_2]) \tag{4.5}$$

As a notational shorthand, we write $f : A \xrightarrow{1\text{-}1} B$. A one-to-one function is also called an *injection*.

Note that the definition of one-to-one in (4.5) is the converse of property F2, which says

$$(\forall x_1, x_2 \in A)([x_1 = x_2] \to [f(x_1) = f(x_2)])$$

Example 4.2.9 Show that the function from Example 4.1.5 is an injection.

Solution Pick $x_1, x_2 \in \mathbb{N}_n$ and suppose that $T(x_1) = T(x_2)$. Then $x_1 + m = x_2 + m$, so that $x_1 = x_2$ by cancellation. Thus T is one-to-one. ■

Exercise 4.2.10 The mapping $f : \mathbb{R} \to \mathbb{R}$ defined by $f(x) = x^3 - 1$ is a function because it is algebraic and its natural domain is the set of all real numbers. Prove that f is one-to-one.

Exercise 4.2.11 Let A be a non-empty set, and $i : A \to A$ the identity function. Show that i is one-to-one.

With Exercises 4.1.8, 4.2.5, and 4.2.11, you have proved the following.

Theorem 4.2.12 For a non-empty set A, the identity mapping is a one-to-one function from A onto itself.

Exercise 4.2.13 What does it mean to say that a function is not one-to-one?

Exercise 4.2.14 Show that $f : [-2, 3] \to [0, 10]$ defined by $f(x) = x^2 + 1$ is not one-to-one.

If $f : A \to B$ is both one-to-one and onto, we may write $f : A \xrightarrow[\text{onto}]{\text{1-1}} B$. A one-to-one, onto function is also called a *bijection*. A bijection from A to B is said to put the sets A and B into a *one-to-one correspondence*.

Exercise 4.2.15 Here are several pairs of sets.

(a) $\{1, 2, 3, 4\}$ and $\{a, b, c\}$

(b) $\{1, 2, 3, 4\}$ and $\mathbb{N}$

(c) $\mathbb{Z}$ and $\mathbb{Z}$

(d) $\mathbb{Z}$ and $\mathbb{W}$

(e) $\mathbb{R}$ and $\mathbb{R}$

For each of the above pairs of sets, find four functions $\{f_1, f_2, f_3, f_4\}$ from the first set to the second set with the following properties, if such functions are possible. It is not necessary to prove that your functions have the desired properties, but be prepared to provide as much explanation as necessary to support your claim.

(i) f_1 is one-to-one but not onto.

(ii) f_2 is onto but not one-to-one.

(iii) f_3 is both one-to-one and onto.

(iv) f_4 is neither one-to-one nor onto.

EXERCISE 4.2.16 In this exercise, we want to consider functions $f : S \to \mathbb{R}$, where f is of the form

$$f(x) = \frac{ax+b}{cx+d} \tag{4.6}$$

where c and d are not *both* zero and where S is the natural domain of f. Note that f is algebraic, so that it has properties F1–F2 on its natural domain.

(a) First we determine S.

  (i) Suppose $c = 0$. By assumption, c and d are not both zero, so $d \neq 0$. What is the natural domain of f?

  (ii) Suppose $c \neq 0$. What is the natural domain of f?

(b) What conditions on a, b, c, and d will guarantee that f is one-to-one? Determine this condition, and then write a complete proof that f, with this additional hypothesis condition, is one-to-one.[2]

(c) From this point forward, assume that the condition from part (b) holds. Thus f is a one-to-one function from S to the real numbers.

  (i) Suppose $c = 0$. Show that f maps onto the real numbers.

  (ii) Suppose $c \neq 0$. Determine the value of y_0 that has no pre-image in S, and prove that f maps S onto $\mathbb{R} - \{y_0\}$.

4.3 Image and Pre-Image Sets

Suppose $f : A \to B$ is a function. Given $A_1 \subseteq A$, we sometimes want to talk about the subset of B that consists of the (unique) images of all the elements of A_1. Also, given $B_1 \subseteq B$, we sometimes talk about the subset of A that consists of all the pre-images of all the elements of B_1. In this section, we define these set constructions and investigate some of their properties.

Definition 4.3.1 Suppose $f : A \to B$ is a function and $A_1 \subseteq A$. Then the *image of* A_1, denoted $f(A_1)$, is defined by

$$f(A_1) = \{\, y \in B : (\exists x \in A_1)(y = f(x)) \,\} \tag{4.7}$$

Thus $y \in f(A_1)$ if and only if there exists $x \in A_1$ such that $y = f(x)$. (See Fig. 4.4.)

[2] If the statement $f(x_1) = f(x_2) \to x_1 = x_2$ is going to be a true statement, what must be assumed about a, b, c, and d to make it so? Play with the statement $f(x_1) = f(x_2)$.

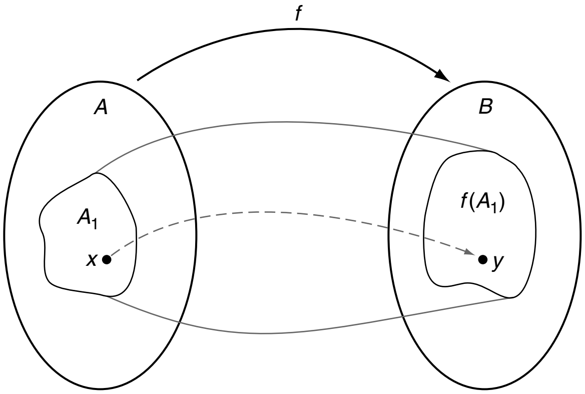

Figure 4.4 The image of A_1 under f from Definition 4.3.1.

With the notation of Definition 4.3.1, we can define the range of a function $f : A \to B$ by

$$\operatorname{Rng} f = f(A) = \{\, y \in B : (\exists x \in A)(y = f(x)) \,\} \tag{4.8}$$

Exercise 4.3.2 Define $f : \mathbb{R} \to \mathbb{R}$ and $g : \mathbb{R} \to \mathbb{R}$ by $f(x) = x^2$ and $g(x) = x^3 - x$. Let $A_1 = \{-2, -1, 0, 1, 2\}$. Construct $f(A_1)$ and $g(A_1)$.

Exercise 4.3.3 Suppose $f : A \to B$ is a function, and let A_1 and A_2 be subsets of A. In the following two set equality statements, three of the four subset inclusion statements are true, and one is false. Prove the three that are true, and provide a counterexample to demonstrate that the fourth is false.

(a) $f(A_1 \cap A_2) = f(A_1) \cap f(A_2)$

(b) $f(A_1 \cup A_2) = f(A_1) \cup f(A_2)$

Exercise 4.3.4 Suppose $f : A \to B$ is a one-to-one function. Show that the false subset inclusion statement from Exercise 4.3.3 is true with this additional assumption.

Exercise 4.3.5 Let $f : A \to B$ be a function and let $\mathcal{F} = \{\, A_\alpha \,\}_{\alpha \in \mathcal{A}}$ be a family of subsets of A. For the three subset inclusion statements in Exercise 4.3.3 that are true, state and prove analogous theorems for $\mathcal{F}$.

Exercise 4.3.6 Suppose $f : A \to B$ is a function and $A_1 \subseteq A_2 \subseteq A$. Then $f(A_1) \subseteq f(A_2)$.

Exercise 4.3.7 Suppose $f : A \to B$ is a function and $A_1 \subseteq A$.

(a) The statement $f(A_1^C) = [f(A_1)]^C$ is not true in general. In fact, neither direction of the subset inclusion is true. Demonstrate these facts by constructing a counterexample to each of the statements $f(A_1^C) \subseteq [f(A_1)]^C$ and $f(A_1^C) \supseteq [f(A_1)]^C$.

(b) Determine an additional hypothesis condition on f that will guarantee $f(A_1^C) \subseteq [f(A_1)]^C$, and prove your claim.

(c) Determine an additional hypothesis condition on f that will guarantee $f(A_1^C) \supseteq [f(A_1)]^C$, and prove your claim.

Now we turn to the pre-image set of a subset of the codomain.

Definition 4.3.8 Given a function $f : A \to B$, and $B_1 \subseteq B$, we define the set $f^{-1}(B_1)$, the *pre-image* of B_1 by

$$f^{-1}(B_1) = \{x \in A : f(x) \in B_1\} \tag{4.9}$$

Thus $x \in f^{-1}(B_1)$ if and only if $f(x) \in B_1$. If $B_1 = \{y\}$ has only one element, we generally write $f^{-1}(y)$ instead of $f^{-1}(\{y\})$. (See Fig. 4.5.)

Exercise 4.3.9 Define $f : \mathbb{R} \to \mathbb{R}$ and $g : \mathbb{R} \to \mathbb{R}$ by $f(x) = x^2$ and $g(x) = x^3 - x$. Let $B_1 = \{-9, 0, 4\}$. Determine $f^{-1}(B_1)$ and $g^{-1}(0)$.

Exercise 4.3.10 Suppose $f : A \to B$ is a function and $B_1 \subseteq B_2 \subseteq B$. Then $f^{-1}(B_1) \subseteq f^{-1}(B_2)$.

Exercise 4.3.11 Let $f : A \to B$ be a function, and suppose B_1 and B_2 are subsets of B. Prove the following set equality statements, if possible. If you claim

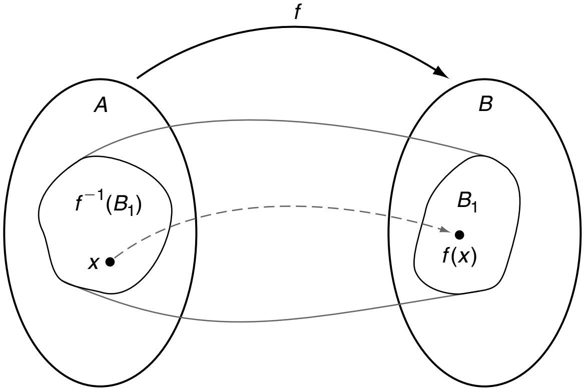

Figure 4.5 The pre-image of B_1 under f from Definition 4.3.8.

that any of the subset inclusion statements are false, provide a counterexample to verify your claim.

(a) $f^{-1}(B_1 \cap B_2) = f^{-1}(B_1) \cap f^{-1}(B_2)$

(b) $f^{-1}(B_1 \cup B_2) = f^{-1}(B_1) \cup f^{-1}(B_2)$

(c) $f^{-1}(B_1^C) = [f^{-1}(B_1)]^C$

EXERCISE 4.3.12 Define $f : \mathbb{R} \to \mathbb{R}$ by $f(x) = x^2$. Let $A_1 = \{-1, 2\}$ and $B_1 = \{-9, 0, 4\}$. Determine $f^{-1}[f(A_1)]$ and $f[f^{-1}(B_1)]$.

EXERCISE 4.3.13 Suppose $f : A \to B$ is a function, $A_1 \subseteq A$ and $B_1 \subseteq B$.

(a) Prove the following.

(i) $A_1 \subseteq f^{-1}[f(A_1)]$

(ii) $f[f^{-1}(B_1)] \subseteq B_1$

(b) Exercise 4.3.12 shows that the reverse subset directions from part (a) do not hold. However, each of these can be proved in the reverse direction with one additional hypothesis condition. Determine the needed additional condition on f for each of the subset statements that will make the reverse subset inclusion statement true. Prove your claims.

EXERCISE 4.3.14 Let $f : A \to B$ be a function, $A_1 \subseteq A$ and $B_1 \subseteq B$.

(a) If $x \in A_1$, what can you say about $f(x)$?

(b) If $y \in f(A_1)$, what does this mean?

(c) If $x \in A$ and $f(x) \in f(A_1)$, then is it necessarily true that $x \in A_1$?

(d) If $x \in f^{-1}[f(A_1)]$, then is it necessarily true that $x \in A_1$?

(e) If $x \in A_1$, then is it necessarily true that $x \in f^{-1}[f(A_1)]$?

(f) If $x \in f^{-1}(B_1)$, then what can you say about $f(x)$?

(g) If $f(x) \in B_1$, then what can you say about x?

(h) If $y \in B_1$, then is it necessarily true that $y \in f[f^{-1}(B_1)]$?

(i) If $y \in f[f^{-1}(B_1)]$, then is it necessarily true that $y \in B_1$?

4.4 Composition and Inverse Functions

A function is a sort of linking of a set A to a set B with certain restrictions. If A is linked to B by a function f, and B is linked to C by a function g, then we can link A to C via B, using f and g in succession. *Composition* is the term we use when combining functions in this way. In this section we will define composition of functions and then investigate the possibility of using $f : A \to B$ to link B back to A by reversing the function f. Such a reverse function, if one exists, is called an *inverse* of f.

4.4.1 Composition of Functions

Suppose $f : A \to B$ and $g : B \to C$ are functions. We want to define a new mapping $g \circ f : A \to C$ that uses the rules of f and g together as in Figure 4.6.

Definition 4.4.1 Given functions $f : A \to B$ and $g : B \to C$, we define the *composition* $g \circ f : A \to C$ to be the mapping defined by $(g \circ f)(x) = g(f(x))$ for all $x \in A$.

Example 4.4.2 Consider the functions $f : \mathbb{R} \to \mathbb{R}^+$ defined by $f(x) = x^2 + 1$ and $g : \mathbb{R}^+ \to \mathbb{R}$ defined by $g(y) = \sqrt{y}$. Then $(g \circ f)(x) = \sqrt{x^2 + 1}$. ■

Notice we did not presume to say that $g \circ f$ is a function in Definition 4.4.1. We must use the fact that f and g are themselves functions to show that properties F1–F2 hold for $g \circ f$.

Exercise 4.4.3 Given functions $f : A \to B$ and $g : B \to C$, the mapping $g \circ f : A \to C$ from Definition 4.4.1 is a function.

We want to address several questions that involve relationships between f, g, and $g \circ f$ based on their individual characteristics. For example, if f and g are one-to-one, does it follow that $g \circ f$ is one-to-one? What about onto? If you know something about f and $g \circ f$, can you conclude anything about g? In the following exercise, we compose several statements of this sort. We call them questions (Q1–Q6) because you must determine whether they are true or false. Not only are some of them true, but some are true even in the absence of one of the given

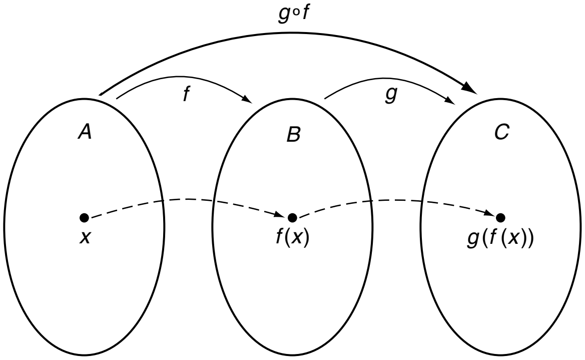

Figure 4.6 Composition: $(g \circ f) : A \to C$.

hypothesis conditions. As an example, we prove that the first of the statements is true.

Exercise 4.4.4 Let $f : A \to B$ and $g : B \to C$ be given functions. Prove each of the following statements, if possible, or provide a counterexample to demonstrate that the statement is false. If a proof does not require one of the provided hypothesis conditions, note how the hypothesis may be relaxed and still be a true statement.

(Q1) If f and g are one-to-one, then $g \circ f$ is one-to-one.

> ***Proof.*** Suppose f and g are one-to-one functions, and suppose $(g \circ f)(x_1) = (g \circ f)(x_2)$. Thus $g[f(x_1)] = g[f(x_2)]$. Since g is one-to-one, we have $f(x_1) = f(x_2)$. Similarly, since f is one-to-one, it follows that $x_1 = x_2$. Thus $g \circ f$ is one-to-one. □

(Q2) If f and $g \circ f$ are one-to-one, then g is one-to-one.

(Q3) If g and $g \circ f$ are one-to-one, then f is one-to-one.

(Q4) If f and g are onto, then $g \circ f$ is onto.

(Q5) If f and $g \circ f$ are onto, then g is onto.

(Q6) If g and $g \circ f$ are onto, then f is onto.

Exercise 4.4.5 Of the Q1–Q6 that are false, one hypothesis condition can be replaced with a different condition to arrive at a true statement. Determine the needed hypothesis condition for each of these, and prove that the resulting statement is true.

4.4.2 Inverse Functions

If $f : A \to B$ is a function, then the rule linking A to B is well defined (F2) on all of A (F1). We now ask a question from the perspective of B. Given $f : A \to B$, can we find a function $g : B \to A$ whose rule is the exact reversal of the rule for f? For example, if $f(8) = -2$, then we want $g(-2) = 8$. Some functions have an inverse, and some do not. First, we establish the criteria by which a function g qualifies as an inverse of f; then we present a theorem that states when such an inverse exists, and that it is unique.

Definition 4.4.6 Suppose $f : A \to B$ is a function. Then we say that a function $g : B \to A$ is an *inverse* of f if $(g \circ f)(x) = x$ for all $x \in A$ and $(f \circ g)(y) = y$ for all $y \in B$. Such a function g is generally denoted f^{-1}.

We must make a careful distinction at this point. In Section 4.3, we used the notation $f^{-1}(B_1)$ to represent the pre-image set of a subset of the codomain of f. This notation has nothing to do with an inverse function f^{-1} per se, and the notation $f^{-1}(B_1)$ can be used for any function f, even if no inverse function for f exists. Furthermore, in Section 4.3 we used the notation $f^{-1}(y)$ to represent the pre-image set of the single element set $\{y\}$, which quite honestly is an abuse of notation. To write $f^{-1}(8)$ appears to be the evaluation of a particular function at a particular point, but it is not necessarily so. If you do not already know that f has an inverse function, then $f^{-1}(8)$ is the subset of the domain of f consisting of all elements that map to 8 under f. However, if a function f^{-1} does exist, then $f^{-1}(8)$ is indeed a unique value.

Here is the theorem that states precisely when a function has an inverse and that such an inverse is unique.

Theorem 4.4.7 Suppose $f : A \to B$ is a one-to-one function from A onto B. Then there exists a unique inverse function $f^{-1} : B \xrightarrow[\text{onto}]{\text{1-1}} A$.

The proof of Theorem 4.4.7 has several parts to it. You will supply all the details in the next exercise, with the help of an outline. Our approach will be to define a new function $g : B \to A$ by appealing to f, then show that g is a one-to-one function from B onto A. Next, by showing that g has the properties required by Definition 4.4.6, g will have earned the right to be called an inverse of f. Finally, you will show that the inverse is unique.

To define the rule for $g : B \to A$, we must pick $y \in B$ and explain how to find some $x \in A$ that we will call $g(y)$. We do this in the following way. Given any $y \in B$, we define $g(y)$ to be any solution x to the equation $f(x) = y$.

Exercise 4.4.8 Complete the following proof of Theorem 4.4.7 by supplying all the missing details.

Proof of Theorem 4.4.7. Let $f : A \xrightarrow[\text{onto}]{\text{1-1}} B$ be a given function. Define a mapping $g : B \to A$ in the following way. For a given $y \in B$, let $g(y)$ be any solution x to the equation $f(x) = y$.

(a) Show that g is defined on all of B(F1).

(b) Show that g is well defined (F2).

(c) Show that g is one-to-one.

(d) Show that g is onto.

(e) Show that $(g \circ f)(x) = x$ for all $x \in A$.

(f) Show that $(f \circ g)(y) = y$ for all $y \in B$.

(g) Finally, suppose g_1 and g_2 are two functions from B to A that satisfy $(g_1 \circ f)(x) = (g_2 \circ f)(x) = x$ for all $x \in A$ and $(f \circ g_1)(y) = (f \circ g_2)(y) = y$ for all $y \in B$. Show that $g_1 = g_2$ by showing $g_1(y) = g_2(y)$ for all $y \in B$.[3] □

[3] All you need are that $(f \circ g_1)(y) = (f \circ g_2)(y) = y$ for all $y \in B$ and that f is one-to-one.

The next exercise brings together many of the results you have derived to this point. Its proof should require little more than an appeal to several results you have already proved.

EXERCISE 4.4.9 For non-empty sets A and B, define $A \equiv B$ provided there exists a one-to-one function from A onto B, that is, there exists a bijection from A to B. Then $\equiv$ defines an equivalence relation on the family of all non-empty sets.

4.5 Three Helpful Theorems

In Sections 4.6 and 4.7 we will discuss *set cardinality*, which is a fancy expression for the number of elements in a set. In Section 4.6 we define what it means for a set to be *finite*, and in Section 4.7 we look at infinite sets. In preparation for the definition of *finite set* and the rigor we require in proving theorems about finite sets, three preliminary theorems will be helpful.

These theorems are not the prettiest ones you will ever see in your mathematical life, but they are not the ugliest either. They are not complicated in principle, but proving them requires slavish attention to some rather minute details. If nothing else, they illustrate the inescapable fact that laying the groundwork for later, more elegant results sometimes requires you to muddle your way through preliminary theorems whose truth is transparent, but whose proofs are a bit sloppy and involve multiple cases. You will prove the first of these theorems. The proof of the second will be provided here, because the sneaky technique it employs will come in handy as you prove the third.

Our first result says that if you are given two one-to-one, onto functions defined from disjoint domains to disjoint codomains, then it is possible to paste together[4] the domains, paste together the codomains, and then define a single function from the pasted domains to the pasted codomains that is also one-to-one and onto.

EXERCISE 4.5.1 Suppose $f_1 : A_1 \to B_1$ and $f_2 : A_2 \to B_2$ are bijections, and that $A_1 \cap A_2 = B_1 \cap B_2 = \emptyset$. Define $f : A_1 \cup A_2 \to B_1 \cup B_2$ by

$$f(x) = \begin{cases} f_1(x), & \text{if } x \in A_1 \\ f_2(x), & \text{if } x \in A_2 \end{cases}$$

Then f is a bijection from $A_1 \cup A_2$ to $B_1 \cup B_2$.

In Section 4.6 we will talk about sets that contain n elements. We place this on firm mathematical footing by using the set $\mathbb{N}_n = \{1, 2, \dots, n\}$ as the domain of a one-to-one function from $\mathbb{N}_n$ onto the given set. The creation of such a function formalizes the way you would count the elements of the set one at a time $1, 2, \dots, n$ and then declare that the set has n elements. It also focuses the entire discussion of all finite sets onto the family of sets of the form $\mathbb{N}_n$ for all positive integers n.

[4] That is, form the disjoint union.

The next helpful theorem will allow us to prove in Section 4.6 that any subset of a finite set is finite. That is, if B has n elements and A is a subset of B, then there is some $m \leq n$ such that A has m elements. We prove this theorem here, for it is a little sticky and will prove to be a good warm-up for proving the third.

Theorem 4.5.2 Let n be a positive integer, and let $S \subseteq \mathbb{N}_n$ be any *non-empty* set. Then there exists a positive integer $m \leq n$ and some $f : S \to \mathbb{N}_m$ where f is a one-to-one, onto function.

Proof. We prove by induction on $n \geq 1$.

(I1) If $n = 1$, then $S = \mathbb{N}_n = \{1\}$. Letting $m = 1$ we may define $f : S \to \mathbb{N}_m$ by $f(1) = 1$, which is clearly a one-to-one, onto function.

(I2) Suppose $n \geq 1$ and that the result is true for any non-empty subset of $\mathbb{N}_n$. Let S be any non-empty subset of $\mathbb{N}_{n+1}$, and consider the following three cases.

If $S = \mathbb{N}_{n+1}$, then we may let $m = n + 1$ and $f : S \to \mathbb{N}_m$ be the identity function.

If $S \subset \mathbb{N}_{n+1}$ and $n + 1 \notin S$, then $S \subseteq \mathbb{N}_n$. Thus the inductive assumption applies, and there exists $m \leq n$ and a function $f : S \xrightarrow[\text{onto}]{\text{1-1}} \mathbb{N}_m$.

If $S \subset \mathbb{N}_{n+1}$ and $n + 1 \in S$, then there exists $k \in \mathbb{N}_n$ such that $k \notin S$. Let $T = [S \cup \{k\}] - \{n + 1\}$. (Draw a picture!) Then $T \subseteq \mathbb{N}_n$, and the inductive assumption applies to yield $m \leq n$ and a function $f : T \xrightarrow[\text{onto}]{\text{1-1}} \mathbb{N}_m$. Define $g : S \to \mathbb{N}_m$ by

$$g(x) = \begin{cases} f(x), & \text{if } x \in T - \{k\} \\ f(k), & \text{if } x = n + 1 \end{cases} \tag{4.10}$$

Defining $A_1 = T - \{k\}$, $A_2 = \{n + 1\}$, $B_1 = \mathbb{N}_m - \{f(k)\}$, and $B_2 = \{f(k)\}$, we may apply Theorem 4.5.1 to conclude that $g : S \to \mathbb{N}_m$ is one-to-one and onto.

□

The last theorem of this section claims that a certain type of one-to-one, onto function does not exist. It will allow us to prove in Section 4.6 that the number of elements in a finite set is well defined. Thus if Nelson says a set has n elements, while Minnie says it has m elements, then the only way they can both be right is if $m = n$. A suggestion for the proof of Exercise 4.5.3 is to use induction on $m \geq 1$.

EXERCISE 4.5.3 Suppose m and n are positive integers with $m < n$. Then there is no bijection $f : \mathbb{N}_n \to \mathbb{N}_m$.[5,6]

4.6 Finite Sets

What does it mean to say that a set has n elements? Or given two sets, what does it mean to say they have the same number of elements? The term that denotes the number of elements in a set is *cardinality*. Consider the following sets.

$$A = \{-5, -2, 1, 4, 7, 10\} \quad \text{and} \quad B = \{\text{b}, \text{d}, \text{h}, \text{p}, \text{f}, \text{l}\} \tag{4.11}$$

What would you say is the cardinality of A? Would you be inclined to say that A and B have the same cardinality? If the cardinality of a set is finite, whatever that might mean, then saying it has cardinality n ought to be a pretty straightforward term to define. Also, if two sets are finite, then saying they have the same cardinality ought to suggest a pretty natural relationship between them.

To determine the cardinality of A defined in Eq. (4.11), you probably counted six elements, then used that as a basis for saying A has cardinality six. Doing the same for B, you then could say that A and B have the same cardinality. To count elements in this way is precisely the motivation behind the following definition.

Definition 4.6.1 Let A be a set, and suppose there exists a positive integer n and a bijection $f : \mathbb{N}_n \to A$. Then we say that A is a *finite* set and that it has *cardinality* n, which we write $|A| = n$. For the empty set, we make a special definition. Let $\mathbb{N}_0$ be defined as the empty set, and define the *empty mapping* $f : \mathbb{N}_0 \to \emptyset$ in order to say that the empty set is finite and has cardinality zero.

To define $f : \mathbb{N}_0 \to \emptyset$ as an empty mapping does not really jibe with Definition 4.1.1, for in our definition of function, we required that domains and codomains be non-empty. But there is no reason we cannot extend the definition of function to include this case, as long as we make sure that it has properties F1–F2. And it does. In fact, the empty mapping meets all the requirements for being a bijection from the empty set to itself, because all the requirements involve the universal quantifier. For example, the empty mapping is onto, for if it were not, then there would exist some $y \in \emptyset$ which has no pre-image. But no such y exists, so the empty mapping is onto. This illustrates an odd truth in

[5] Prove the inductive step by contrapositive.

[6] If $f : \mathbb{N}_{n+1} \to \mathbb{N}_{m+1}$ is one-to-one and onto, and if $f(n+1) = m+1$, then defining g should be easy. Otherwise, $f(n+1) = k$ for some $1 \le k \le m$ and $f(\ell) = m+1$ for some $1 \le \ell \le n$. Define g to map ℓ to k.

mathematics: Statements of the form $(\forall x \in \emptyset)(P(x))$ are true, regardless of $P(x)$. For example, every negative natural number is a proper divisor of 11.

Example 4.6.2 Show that the set B in (4.11) has cardinality 6.

Solution Define $f : \mathbb{N}_6 \to B$ in the following way. Let $f(1) = \text{b}$, $f(2) = \text{d}$, $f(3) = \text{h}$, $f(4) = \text{p}$, $f(5) = \text{f}$, $f(6) = \text{l}$. Since f is one-to-one and onto B, we have shown $|B| = 6$. ■

Example 4.6.3 By the identity mapping, $|\mathbb{N}_n| = n$. ■

Definition 4.6.1 introduces two terms concerning the size of some but not all sets: *finiteness* and *cardinality*. To say a set A is finite is to mean that it has some nonnegative integer n associated with it, called its cardinality, which derives from the existence of some bijection from $\mathbb{N}_n$ to A. Conversely, to say that $|A| = n$, where n is a nonnegative integer, is to say that A is finite. Notice that the empty set is the only set with cardinality zero. For if A is any non-empty set, then no function from $\mathbb{N}_0$ to A can be onto.

With Example 4.6.2, we have earned the right to make a statement such as "B has six elements," and we have validated our inclination to determine its cardinality by walking our fingers over its elements and counting them. We must be careful, though. The existence of a bijection from $\mathbb{N}_n$ to A is a basis for declaring the cardinality of A to be n. But how do we know that there is not some different m and a bijection from $\mathbb{N}_m$ to A? If there were such an m, then cardinality would not be well defined, for $|A|$ could be two different numbers. With the help of Exercise 4.5.3, your proof of the next result should be quick.

EXERCISE 4.6.4 The cardinality of a finite set A is well defined. That is, if $|A| = m$ and $|A| = n$, then $m = n$.[7]

Exercise 4.4.9 allows us to say that the family of all sets is partitioned into equivalence classes based on the existence of bijections between them. Now we can see what some of these equivalence classes are. Definition 4.6.1 says that the cardinality of some sets is a nonnegative integer n, provided the set is in the equivalence class of $\mathbb{N}_n$. Furthermore, by Exercise 4.6.4, that n is unique. So if $m \neq n$, then $\mathbb{N}_m$ and $\mathbb{N}_n$ are representative elements of different equivalence classes. With this, we have proved the following.

Theorem 4.6.5 Suppose A is a finite set, and B is any set. Then A and B have the same cardinality if and only if there exists a bijection from A to B.

Corollary 4.6.6 If $f : A \to B$ is a one-to-one function, then $|A| = \left|\operatorname{Rng} f\right|$.

[7] Address the empty set as a special case. Otherwise, if the result is not true, it violates Exercise 4.5.3.

The next exercise says that if A and B are disjoint sets, and if there exist bijections $f_1 : \mathbb{N}_m \to A$ and $f_2 : \mathbb{N}_n \to B$, then you can construct a bijection $f : \mathbb{N}_{m+n} \to A \dot{\cup} B$. Once you have defined such a function, showing it is a bijection from $\mathbb{N}_{m+n}$ to $A \dot{\cup} B$ should amount to little more than calling on previous exercises.

EXERCISE 4.6.7 If A and B are disjoint finite sets with cardinalities m and n, respectively, then $|A \cup B| = m + n$.[8]

EXERCISE 4.6.8 If $|B| = n$ and $A \subseteq B$, then A is finite and satisfies $|A| \le n$.[9]

Corollary 4.6.9 Suppose $|A| = n$ and C is any set. Then $|A \cap C| \le |A|$.

Proof. Since $A \cap C \subseteq A$, the result is immediate from Exercise 4.6.8. □

EXERCISE 4.6.10 Suppose U is a finite universal set, and $A \subseteq U$. Then A and A^C are both finite, and $|A| + \left|A^C\right| = |U|$.

EXERCISE 4.6.11 Suppose $|A| = |B| = n$, and suppose $f : A \to B$ is any one-to-one function. Then f is onto.[10]

EXERCISE 4.6.12 If $|A| = m$ and $|B| = n$, then $|A \cup B| \le m + n$.[11]

EXERCISE 4.6.13 The union of a finite number of finite sets is finite.[12]

4.7 Infinite Sets

The following definition should not be too surprising.

Definition 4.7.1 A set is said to be *infinite* provided it is not finite.

By negating Definition 4.6.1, we see that a set A is infinite provided that for every nonnegative integer n, every function from $\mathbb{N}_n$ to A fails to be either one-to-one or onto. In reality, if there were some $f : \mathbb{N}_n \to A$ that is onto but not one-to-one, we could create a function $f_1 : \mathbb{N}_m \to A$ for some $m < n$ that culls out

[8] Make use of Exercise 4.5.1 where the domain of f is $\mathbb{N}_{m+n}$ and the range is $A \cup B$. The function $T : \mathbb{N}_n \to \mathbb{N}_n^m$ from Example 4.1.5 should be helpful, too.

[9] If $f : \mathbb{N}_n \to B$ is a bijection, apply Theorem 4.5.2 to $f^{-1}(A)$. Construct $g : \mathbb{N}_m \xrightarrow{\text{1-1}} A$ by composition.

[10] What if f is not onto?

[11] See Exercise 3.2.14.

[12] Use induction and Exercise 4.6.12.

the repetition in f and is therefore one-to-one.[13] By creating f_1 in this way, we would have that A is finite. The point is stated in the following.

Theorem 4.7.2 A set A is infinite if and only if for every nonnegative integer n and every function $f : \mathbb{N}_n \to A$, f is not onto.

Loosely speaking, if a set is infinite, then no matter how large n is, there are not enough elements in $\mathbb{N}_n$ to tag all the elements of the set.

Before we get into the interesting results of infinite sets, let's point out where it will lead. Strangely, just because two sets are both infinite, it does not follow that they have the same cardinality, in the sense that there is a one-to-one correspondence between them. Some infinite sets are actually bigger than others. This makes for some real surprises and motivates us to discuss different *orders* of infinity. In fact, it is possible to generate an infinite sequence of infinite sets $A_1, A_2, \ldots$ where $|A_n|$ is a higher order of infinity than $|A_{n-1}|$. It's mind boggling. Not only is there more than one size of infinity, but also there are infinitely many infinities. A natural question then arises: Of the infinitely many different orders of infinity, which one represents the number of distinct infinities that exist? We will look at only two sizes of infinity. Here is our first.

Definition 4.7.3 Suppose A is an infinite set. Then A is said to be *countably infinite* provided there exists a one-to-one function from the positive integers $\mathbb{N}$ onto A, and we say that A has cardinality $\aleph_0$ ($\aleph$ is the Hebrew letter *aleph*). If A is finite or countably infinite, we say that A is *countable*. If A is not countable, we say that it is *uncountable*.

Example 4.7.4 By the identity function, the natural numbers are countably infinite. ■

EXERCISE 4.7.5 Show that the whole numbers are countably infinite.

The equivalence relation from Exercise 4.4.9 allows us to say that the natural numbers are a representative set from the equivalence class of all countably infinite sets. Let's see some other sets that are countably infinite and prove some theorems about countable sets in general. Remember, to show a set A is countably infinite, your task is to construct a one-to-one function from the positive integers onto A. Sometimes an easy way to build such a function is by systematically declaring the values of $f(1)$, $f(2)$, $f(3)$, and so on, which is the same as ordering the elements

[13] We begin by creating a subset $S \subseteq \mathbb{N}_n$ so that $f : S \to A$ is one-to-one by applying the axiom of choice (discussed on page 118). For each $a \in A$, we could take a single element of $f^{-1}(a)$, and then let S consist of all these chosen pre-images. Then we could apply Theorem 4.5.2 to conclude the existence of a bijection $g : S \to \mathbb{N}_m$ for some $m < n$. Then $f_1 = f \circ g^{-1}$ would be a one-to-one function from $\mathbb{N}_m$ to A.

of A. This ordering will correspond to a one-to-one, onto function provided every element of A appears in the ordering exactly once. Once you have done this, you might be able to concoct an explicit formula for $f(n)$. In the next exercise, you might find a way to define the function by

$$f(n) = \begin{cases} ? & \text{if } n \text{ is even} \\ ? & \text{if } n \text{ is odd} \end{cases}$$

then apply Exercise 4.5.1.

Exercise 4.7.6 The integers are countably infinite.

If we expect to stumble eventually onto an uncountable set, the rationals might seem to be the first one we would find. After all, in Exercise 2.2.9, you showed that $a < (a + b)/2 < b$, so that between any two rational numbers there is another rational number. Thus the rational numbers are strewn densely up and down the real number line, as opposed to the integers, which have plenty of room in between each one. Well, the rationals are countable. In Theorem 4.7.7 we prove that the positive rational numbers are countable. Although it might sound counterintuitive, it is possible to list the positive rationals sequentially as $f(1)$, $f(2)$, $f(3)$, and so on, to create a function that demonstrates countability.

Theorem 4.7.7 The positive rationals are countable.

Proof. Consider Figure 4.7, which lists all the positive rational numbers. From this figure, we may define f in the following way. Starting in the upper left corner of

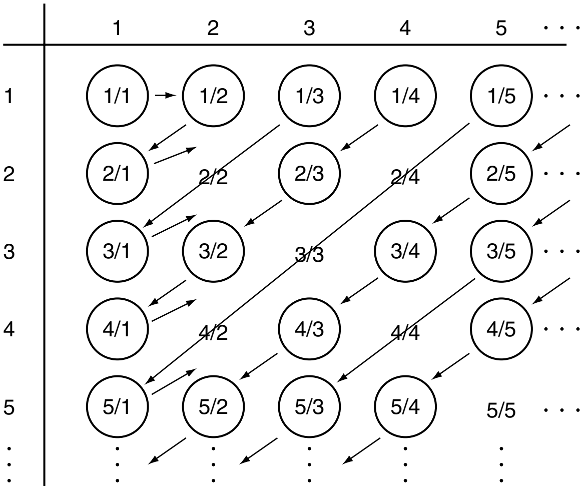

Figure 4.7 Ordering of $\mathbb{Q}^+$ to show countability.

the table, define $f(1) = 1/1$. We then move down successive diagonals from upper right to lower left, defining $f(2) = 1/2$ and $f(3) = 2/1$ from the first diagonal. Moving down the next diagonal, we define $f(4) = 1/3$. But since $2/2 = 1/1$, we skip $2/2$ and define $f(5) = 3/1$. Continuing in this fashion,

$$f(6) = \frac{1}{4}, \quad f(7) = \frac{2}{3}, \quad f(8) = \frac{3}{2}, \quad f(9) = \frac{4}{1}, \quad f(10) = \frac{1}{5}, \quad f(11) = \frac{5}{1}$$

and so on, making sure f is one-to-one by skipping over rational numbers that are not in reduced form. This program guarantees that f is one-to-one and onto, so we have shown that the positive rationals are countable. □

Now let's pause for a moment and reflect on the proof of Theorem 4.7.7. We can think of constructing the one-to-one correspondence as sequencing the elements of A as $\langle a_1, a_2, a_3, \ldots \rangle$, where every element of A is in the list exactly once. If such a listing of elements of A can be found, then you have shown that A is countably infinite.

In tracing through the grid in the proof of Theorem 4.7.7, we ensured that f is one-to-one by skipping over entries that were not in reduced form. If we had not done this skipping, the resulting f would still have been onto, but not one-to-one. Knowing a priori that the positive rationals are infinite, this failure of f to be one-to-one would not be particularly disconcerting. Loosely speaking, here's why. If there are enough positive integers to tag all the positive rational numbers with some repetition, then there ought to be enough to tag them without repetition. The next theorem gives us the freedom not to worry about this repetition and will save us from some minor headaches later. It deals with all countable sets, whether they are finite or countably infinite. We state it here without proof.

Theorem 4.7.8 A non-empty set A is countable if and only if there exists a function from the positive integers onto A.

The convenience of Theorem 4.7.8 will become apparent in the following results, where we need not worry if a given countable set is finite or countably infinite. For example, in Exercise 4.7.9, given two countable sets A and B, Theorem 4.7.8 will allow you to prove that $A \cup B$ is countable by constructing a mapping from the positive integers onto $A \cup B$ without having to worry whether A and B are disjoint or whether they are finite or countably infinite. Since the union across a family of sets is unchanged if all empty sets in the family are omitted, we are free to assume that all sets in the following exercises are non-empty.

Exercise 4.7.9 If A and B are countable, then $A \cup B$ is countable.

Exercise 4.7.10 If $\{A_k\}_{k=1}^n$ is a finite family of countable sets, then $\cup_{k=1}^n A_k$ is countable.

EXERCISE 4.7.11 If $\{A_n\}_{n=1}^{\infty}$ is a countably infinite family of countable sets, then $\cup_{n=1}^{\infty} A_n$ is countable.[14]

By Exercise 4.7.10, it follows that the rationals are countable, for $\mathbb{Q} = \mathbb{Q}^+ \cup \mathbb{Q}^- \cup \{0\}$, and $\mathbb{Q}^-$ is countable in the same way that $\mathbb{Q}^+$ is.

If every function $f : \mathbb{N} \to A$ fails to be onto, then A is uncountable. This is how we prove the following.

Theorem 4.7.12 $[0, 1]$ is uncountable.

Proof. Suppose $f : \mathbb{N} \to [0, 1]$ is *any* function. We can show that f fails to be onto. From assumption A21 in Chapter 0, we may represent $[0, 1]$ as the set of all decimal expansions of the form 0.XXXXX.... Consider the following diagram.

$$\begin{aligned}
f(1) &= 0\,.\,a_1\,a_2\,a_3\,a_4\ldots\\
f(2) &= 0\,.\,b_1\,b_2\,b_3\,b_4\ldots\\
f(3) &= 0\,.\,c_1\,c_2\,c_3\,c_4\ldots\\
f(4) &= 0\,.\,d_1\,d_2\,d_3\,d_4\ldots\\
f(5) &= 0\,.\,e_1\,e_2\,e_3\,e_4\ldots\\
&\vdots
\end{aligned}$$

We can show f is not onto by constructing a real number in $[0, 1]$ that is not in the list. Let $x = 0.x_1x_2x_3x_4\ldots$ where $x_1 \neq a_1$, $x_2 \neq b_2$, $x_3 \neq c_3$, etc.; and no $x_k = 9$.[15] Then $x \in [0, 1]$, and x is different from every number in the range of f. Thus $[0, 1]$ is not countable. □

EXERCISE 4.7.13 If B is countable and $A \subseteq B$, then A is countable.[16]

Corollary 4.7.14 The real numbers are uncountable.

Proof. Since $[0, 1]$ is uncountable and $[0, 1] \subset \mathbb{R}$, the result follows from Exercise 4.7.13. □

EXERCISE 4.7.15 The irrationals are uncountable.

[14] Induction won't help here. Try a grid like that in the proof of Theorem 4.7.7.

[15] This way we guarantee that x does not settle into the pattern of the repeating 9.

[16] Use Theorem 4.7.8 to take care of the case where$A$ is non-empty.

4.8 Cartesian Products and Cardinality

Combinatorics is the area of mathematics that deals with techniques for counting. In this section and the next, we study some useful counting techniques whose mathematical foundation lies in the cardinality of Cartesian products and the number of functions from one finite set to another.

To simplify some of the language and notation, if A is a finite set and f is a one-to-one function from $\mathbb{N}_n$ onto A, we will address elements of A as $\{a_1, a_2, \dots, a_n\}$, where $a_k = f(k)$ for $1 \le k \le n$. Since f is a function, every k can be used to reference a unique $a_k \in A$. Since f is one-to-one and onto, every element of A can be addressed as a_k for some unique k.

4.8.1 Cartesian Products

In Section 3.11, we defined the *Cartesian product* of two sets A and B as

$$A \times B = \{(a, b) : a \in A, b \in B\} \tag{4.12}$$

the set of ordered pairs whose first coordinate comes from A and whose second coordinate comes from B. Assuming A and B each have a form of equality defined on them, which we temporarily write as $=_A$ and $=_B$, we can define equality in $A \times B$ by

$$(a_1, b_1) =_{A\times B} (a_2, b_2) \quad \text{if and only if} \quad a_1 =_A a_2 \text{ and } b_1 =_B b_2 \tag{4.13}$$

Exercise 4.8.1 Assuming $=_A$ and $=_B$ are equivalence relations on A and B, respectively, show that $=_{A\times B}$ defines an equivalence relation on $A \times B$.

From this point forward, we will simply write that $(a_1, b_1) = (a_2, b_2)$ provided $a_1 = a_2$ and $b_1 = b_2$. Our first theorem will serve you as the root of an induction argument for Exercise 4.8.5.

Theorem 4.8.2 Suppose $|A| = 1$ and $|B| = n \ge 1$. Then $|B| = |A \times B|$.

We prove Theorem 4.8.2 by constructing a bijection from B to $A \times B$, thereby demonstrating that B and $A \times B$ are in the same cardinality class.

Proof. Write $A = \{a\}$ and $B = \{b_1, b_2, \dots b_n\}$. Define $f : B \to A \times B$ by $f(b_k) = (a, b_k)$. First, for any $b_k \in B$, $(a, b_k) \in A \times B$, so f has property F1. Also, if $b_j = b_k$, then $f(b_j) = (a, b_j) = (a, b_k) = f(b_k)$, so that f has property F2. If we suppose $f(b_j) = f(b_k)$, then $(a, b_j) = (a, b_k)$, which implies $b_j = b_k$, and f is one-to-one. Finally, if we pick any $(a, b_k) \in A \times B$, then $b_k \in B$ and $f(b_k) = (a, b_k)$, so that f is onto. Since we have found a one-to-one correspondence between B and $A \times B$, they are in the same cardinality class, and $|B| = |A \times B|$. □

Exercise 4.8.3 If A_1, A_2, and B are sets, then

$$(A_1 \cup A_2) \times B = (A_1 \times B) \cup (A_2 \times B) \tag{4.14}$$

Furthermore, if A_1 and A_2 are disjoint, then so are $(A_1 \times B)$ and $(A_2 \times B)$.

Corollary 4.8.4 If A_1, A_2, and B are finite sets, where A_1 and A_2 are disjoint, then $|(A_1 \cup A_2) \times B| = |A_1 \times B| + |A_2 \times B|$.

Proof. By Exercises 4.6.7 and 4.8.3,

$$\left|(A_1 \dot{\cup} A_2) \times B\right| = \left|(A_1 \times B) \dot{\cup} (A_2 \times B)\right| = |A_1 \times B| + |A_2 \times B| \tag{4.15}$$

□

You now have all you need to write an induction argument for the following.

Exercise 4.8.5 Suppose A and B are non-empty, finite sets. Then

$$|A \times B| = |A| \times |B|^{17} \tag{4.16}$$

We can form the Cartesian product of more than two sets, but to do it conveniently we need to be a bit less than rigorous. We want $A \times B \times C$ to be the set of ordered triples (a, b, c) where $a \in A$, $b \in B$, and $c \in C$. But to define $A \times B \times C$ in terms of the *binary* Cartesian product defined in Eq. (4.12), we need to associate either A and B, or B and C. This leaves us with $(A \times B) \times C$ or $A \times (B \times C)$, which, unfortunately, are different. As we defined the Cartesian product in Eq. (4.12),

$$(A \times B) \times C = \{((a, b), c) : a \in A, b \in B, c \in C\} \tag{4.17}$$

but

$$A \times (B \times C) = \{(a, (b, c)) : a \in A, b \in B, c \in C\} \tag{4.18}$$

Expressions of the form $((a, b), c)$ and $(a, (b, c))$ are not the same, though it does not create an insurmountable obstacle. Rather than try to deal with the lack of associativity in the Cartesian product of sets, we make the following recursive definition and then give ourselves the freedom to ignore the associativity question.

[17] Induct on $|A| \geq 1$, thinking of $|B|$ as fixed.

Definition 4.8.6 Suppose $\{A_n\}$ is a family of sets indexed by the positive integers. Then we define the sequence of Cartesian products recursively as

$$\prod_{k=1}^{1} A_k = A_1 \tag{4.19}$$

$$\prod_{k=1}^{n+1} A_k = \big(\prod_{k=1}^{n} A_k\big) \times A_{n+1} \quad \text{for } n \geq 1 \tag{4.20}$$

With Definition 4.8.6, elements of $\prod_{k=1}^{n}\{A_k\}$ for the first few values of n would look like this:

$$\begin{aligned}
\prod_{k=1}^{1} A_k &= \{a : a \in A_1\} \\
\prod_{k=1}^{2} A_k &= \{(a_1, a_2) : a_k \in A_k \text{ for } 1 \leq k \leq 2\} \\
\prod_{k=1}^{3} A_k &= \{((a_1, a_2), a_3) : a_k \in A_k \text{ for } 1 \leq k \leq 3\} \\
\prod_{k=1}^{4} A_k &= \{(((a_1, a_2), a_3), a_4) : a_k \in A_k \text{ for } 1 \leq k \leq 4\}
\end{aligned} \tag{4.21}$$

Instead of writing elements of $\prod_{k=1}^{n} A_k$ this way, we give ourselves the freedom to address elements as $(a_1, a_2, \ldots, a_n)$, n-tuples where $a_k \in A_k$ for all $1 \leq k \leq n$. With this definition and another induction argument, you can prove the following.

Exercise 4.8.7 Suppose $\{A_k\}_{k=1}^{n}$ is a finite family of finite sets. Then

$$\left|\prod_{k=1}^{n} A_k\right| = \prod_{k=1}^{n} |A_k| \tag{4.22}$$

4.8.2 Functions Between Finite Sets

Suppose A and B are non-empty sets with $|A| = m$ and $|B| = n$. We can now calculate the number of functions from A to B in the following way. Form the Cartesian product of m copies of B, which we may write B^m. By Exercise 4.8.7, $|B^m| = n^m$. Furthermore, any ordered m-tuple from B^m can be thought of as a function $f : A \to B$, where the first coordinate of the m-tuple is $f(a_1)$, the second coordinate is $f(a_2)$, and so on. Since distinct m-tuples from B^m represent distinct functions from A to B, we have proved the following.

Theorem 4.8.8 If A and B are non-empty sets with $|A| = m$ and $|B| = n$, then there are n^m distinct functions from A to B.

Here is an informal way to think of Theorem 4.8.8. Let $A = \mathbb{N}_m$ and imagine we have m empty slots numbered $1, 2, \ldots, m$. Into each slot we insert precisely one element of B, where any element of B may be used repeatedly if we like. Filling the slots in this way is equivalent to defining a function $f : \mathbb{N}_m \to B$. There are n choices of element to place in the first slot. For each of these ways to fill the first slot, there are n ways to fill the second slot, so that there are n^2 ways to fill the first two slots. Continuing, we have that there are n^m ways to fill all m slots. Theorem 4.8.8 motivates the notation B^A to represent the set of all functions from A to B.

What if we must fill each of the n slots with one of the m elements of B, but we are not allowed to use any element of B more than once? How many ways can this be done? First note that if such a task can be done at all, it must be that $m \leq n$. We have a choice of n objects for the first slot. Then, regardless of which element of B was placed in the first slot, there are $n - 1$ possible elements for the second slot, and so on. Thus the number of ways to fill the m slots with *distinct* elements of B is $n(n-1)(n-2)\ldots(n-m+1)$. We write this as

$$P(n, m) = n(n-1)(n-2)\ldots(n-m+1) = \frac{n!}{(n-m)!} \tag{4.23}$$

which is called the number of *permutations* of n objects taken m at a time. Filling the slots with elements of B so that there is no repetition is the same as constructing a function from B to $\mathbb{N}_m$ that is one-to-one. With this, we have an informal proof of the following. A rigorous proof should be done by induction.

Exercise 4.8.9 Let A and B be non-empty sets with $|A| = m \leq n = |B|$. Then there are $P(n, m)$ distinct one-to-one functions from A to B.[18]

If Eq. (4.23) is to be taken as meaningful for $m = 0$, it implies that there is one way to arrange zero of n objects. It also suggests that there exists a unique (empty) function from the empty set to any finite set.

If U is a finite univeral set and A is a subset of U, then Exercise 4.6.10 implies the *complement rule*, which says

$$|A| = |U| - \left|A^C\right| \tag{4.24}$$

Sometimes the task of determining the cardinality of a set is more easily done by exploiting Eq. (4.24).

[18] Induct on $m \geq 1$.

EXERCISE 4.8.10 Let A and B be non-empty sets with $|A| = m \leq n = |B|$. How many distinct functions from A to B are not one-to-one?

Suppose A is a set with n elements and $(a_1, \ldots, a_n)$ is an element of A^n, where all coordinates of the n-tuple are distinct. We may think of such an n-tuple as a way of ordering the elements of A.

EXERCISE 4.8.11 If A is a set with n elements, in how many distinct ways can its elements be ordered?

4.8.3 Applications

How many ways are there to put together a meal from all the cafeteria offerings? To arrange your CD collection on your dresser? To name a baby boy from a list of family preferences? Results from this section can help us answer some questions of the sort "How many ways are there to ...?"

Example 4.8.12 Your university cafeteria has the following menu for today's lunch.

Meats :	Meatloaf, Chicken, Fish sticks
Starchy vegetables :	Potatoes, Rice, Corn, Pasta
Green vegetables :	Beans, Broccoli, Salad, Spinach
Breads :	Rolls, Corn bread
Desserts :	Chocolate cake, Pudding

If you choose one item from each category of the menu, how many different meals could you put together?

Solution If we let A_1 be the set of meat offerings, A_2 the set of starchy vegetables, and so on, then each potential complete meal is an element of $\prod_{k=1}^{5} A_k$. By Exercise 4.8.7, there are $3 \times 4 \times 4 \times 2 \times 2 = 192$ possible meals. ■

In Example 4.8.12, counting the number of possible meals can be visualized in the following way. We have five empty slots to fill on our plate, the first to be filled with a choice of meat, the second with a choice of starchy vegetable, and so on. Furthermore, the number of choices available to fill a particular slot is unaffected by the way any of the previous slots have been filled or the way the remaining slots will be filled. Multiplying the number of ways to fill each slot illustrates the *multiplication rule*. If an n-step process is such that the kth step can be done in a_k ways, and the number of ways each step can be done is unaffected by the choice made for any other step, then the total number of ways to perform all n steps is $\prod_{k=1}^{n} a_k$.

EXERCISE 4.8.13 How many meals can be put together from the menu in Example 4.8.12 if dessert is optional?

Here is an illustration of the practical use of Theorem 4.8.8.

Example 4.8.14 Suppose you toss a coin 10 times and observe the sequence of outcomes of heads (H) or tails (T). Each possible outcome of the 10 tossings can be written as a 10-tuple of the form (H,H,H,T,H,H,T,T,H,T), each of which we can visualize as a function from $\mathbb{N}_{10}$ to $\{H, T\}$. By Theorem 4.8.8, there are $2^{10} = 1024$ such functions. ■

EXERCISE 4.8.15 A password to your computer account must be precisely seven alphanumeric characters, that is, a sequence of seven characters taken from the 26 letters of the alphabet and the digits 0–9. How many distinct passwords may be formed?

Exercise 4.8.9 will help you answer the next question.

EXERCISE 4.8.16 The Fitzpatricks are expecting a baby boy any day now. Family members have strong opinions about what the child will be named. Suggested names are William, Warren, Benjamin, Fitzhugh, Chancellor, Millhouse, and Nebuchadnezzar. The parents want their son to have three distinct given names, such as William Fitzhugh Millhouse Fitzpatrick. How many potential names are there for the Fitzpatrick son?

EXERCISE 4.8.17 How many ways are there to arrange your collection of eight CDs in a row on top of your dresser?

EXERCISE 4.8.18 In California, a standard license plate consists of a digit 1–9, followed by three letters, followed by three more digits 0–9, such as 3AAG045. If all constructions of this form are considered usable, how many standard license plates can California issue?

EXERCISE 4.8.19 Twenty horses run a race in which prizes are given for win, place, and show (first, second, and third places). How many outcomes are there for the race?

Example 4.8.20 Let a *word* be defined as any arrangement of letters from the alphabet, so that, for example EIEIO is a word. If the vowels are {A, E, I, O, U}, how many four-letter words contain at least one vowel?

Solution Let U be the set of all four-letter words, and let A be the set of four-letter words that contain at least one vowel. Then $|U| = 26^4$ and $\left|A^C\right| = 21^4$. By the complement rule, $|A| = 26^4 - 21^4$. ■

EXERCISE 4.8.21 For security purposes, a password like that in Exercise 4.8.15 must contain at least one numeric character. How many such passwords are there?

The remaining exercises will require you to apply all the principles of this section.

EXERCISE 4.8.22 A single digit on an old-fashioned calculator is created from a configuration of seven light-emitting diodes, as in the sketch here. How many distinct symbols can be displayed with these seven lights?

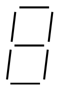

EXERCISE 4.8.23 In a first grade class with 20 students, every student sends a Valentine's Day card to every other student. How many Valentine's Day cards are sent among the students?

EXERCISE 4.8.24 You live in Atlanta and must fly for work to Washington DC, Chicago, and Dallas, then back to Atlanta. If you may visit these cities in any order, how many different ways can you plan your trip?

EXERCISE 4.8.25 Six houses in a row are each to be painted with one of the colors red, blue, green, and yellow. In how many ways may the houses be painted so that no two adjacent houses are the same color?

EXERCISE 4.8.26 The ballot for your club's officers lists 2 candidates for president, 4 candidates for vice-president, and 6 candidates for secretary. In how many different ways can you mark your ballot if:

(a) You vote for exactly one person for each office?

(b) You may choose to leave any item blank?

EXERCISE 4.8.27 The postal delivery person has five different packages, all to be delivered to the same block, where there are 10 houses. How many different ways can he or she distribute the packages among the houses, provided:

(a) There are no restrictions on how the packages may be distributed?

(b) No house may receive more than one package?

EXERCISE 4.8.28 Six friends go out for dinner, and each orders something different. The server brings 4 of the 6 dishes to the table, but he does not remember who ordered what. In how many different ways could he set these 4 dishes down among the 6 people at the table?

4.9 Combinations and Partitions

In this section, we build on this result to determine the number of ways we may select m out of n objects, where we do not distinguish between different arrangements. Then we generalize this to determine how many ways we may partition a given set into distinguishable subsets of given sizes.

4.9.1 Combinations

Given a set of n objects and an integer $0 \leq m \leq n$, we can calculate the number of ways we may choose m of the n objects, with no order taken into consideration. This is precisely the same as asking how many distinct subsets containing m objects may be formed from a set with n objects. We call this the number of *combinations* of n objects taken m at a time, and we denote it $C(n, m)$. An alternate notation is $\binom{n}{m}$, which is read "n choose m."

A natural way to derive the formula for $\binom{n}{m}$ is with an equivalence relation. Let S be the set of all the $P(n, m)$ permutations of the n objects taken m at a time. Define two permutations in S to be equivalent provided they contain precisely the same elements of the set, regardless of the order of these elements. Clearly, this is an equivalence relation, so that S is partitioned into equivalence classes, where any given class contains all the possible arrangements of a particular m-element subset of the n objects. The number of equivalence classes is therefore $C(n, m)$, the expression whose formula we are trying to determine. Notice that each equivalence class contains precisely $m!$ elements. Since $|S| = P(n, m)$, we have that $m! \times C(n, m) = P(n, m)$, so that

$$C(n, m) = \binom{n}{m} = \frac{P(n, m)}{m!} = \frac{n!}{m!(n-m)!} \tag{4.25}$$

We have therefore proved the following.

Theorem 4.9.1 Suppose A is a finite set of cardinality n, and let $0 \leq m \leq n$. Then the number of subsets of A of cardinality m is given by Eq. (4.25).

Notice that $C(n, m) = C(n, n - m)$, so the number of ways to choose m out of n objects is the same as the number of ways to choose $n - m$ of the objects. Notice also that $C(n, 0) = 1$ is meaningful in that there is precisely one zero-element subset of a given set.

Combinations have a number of interesting mathematical features, and they pop up in a variety of mathematical settings. One useful result is the following.

Exercise 4.9.2 Suppose $n \geq 1$ and $1 \leq k \leq n$. Then

$$\binom{n}{k-1} + \binom{n}{k} = \binom{n+1}{k}^{19} \tag{4.26}$$

It is arguable that people who work in combinatorics enjoy their proofs as much as any group of mathematicians. The reason is that there are often several seemingly unrelated ways to argue the same result, and to discover a new way to prove a combinatorial theorem is a particularly satisfying accomplishment. One such *combinatorial argument* for Exercise 4.9.2 goes like this.

Proof of Exercise 4.9.2. Let n be a nonnegative integer, and suppose A is a set of $n+1$ objects. Then the number of k-element subsets of A is given by the right-hand side of Eq. (4.26). Now let a be any element of A. The k-element subsets of A are of two types—those that contain a and those that do not. To create a k-element subset that contains a, we need only choose $k-1$ of the n objects other than a. There are $\binom{n}{k-1}$ ways to do this. To create a k-element subset that does not contain a, we must choose k of the n objects other than a. There are $\binom{n}{k}$ ways to do this. Thus another way to count the number of k-element subsets of A is given by the left-hand side of Eq. (4.26). □

4.9.2 Partitioning a Set

Choosing an m-element subset of an n-element set is the same as partitioning the set into two subsets, one of size m and one of size $n-m$. We can generalize this process in the following way.

Suppose A has n elements and we want to partition A into p subsets whose cardinalities are $n_1, n_2, \ldots, n_p$, respectively, where $\sum_{k=1}^{p} n_k = n$ and none of the n_k is zero. How many ways can we do it? Imagine choosing n_1 of the n objects to comprise the first subset, then choosing n_2 of the remaining $n - n_1$ objects to comprise the second subset, then n_3 of the remaining $n - n_1 - n_2$, and so on. By the multiplication rule, then we have

$$\begin{aligned}
&\binom{n}{n_1}\binom{n-n_1}{n_2}\binom{n-n_1-n_2}{n_3}\cdots\binom{n_p}{n_p} \\
&= \frac{n!}{n_1!(n-n_1)!} \cdot \frac{(n-n_1)!}{n_2!(n-n_1-n_2)!} \cdot \frac{(n-n_1-n_2)!}{n_3!(n-n_1-n_2-n_3)!} \cdot \frac{n_p!}{n_p!0!} \\
&= \frac{n!}{n_1!n_2!n_3! \cdot n_p!}
\end{aligned} \tag{4.27}$$

[19] Induction is not necessary, just some straightforward algebraic manipulation.

Theorem 4.9.3 Suppose $|A| = n$, and let $n_1, \ldots, n_p$ be positive integers such that $\sum_{k=1}^{p} n_p = n$. Then the number of ways to partition A into distinguishable subsets of sizes $n_1, \ldots, n_p$ is given by

$$\binom{n}{n_1, \ldots, n_p} = \frac{n!}{n_1!n_2!\ldots n_p!} \tag{4.28}$$

Theorem 4.9.3 says that the subsets of A are distinguishable. If all of the subset sizes $n_1, \ldots, n_p$ are distinct, then all of the p subsets of A are clearly distinguishable. However, if $n_i = n_j$ for some distinct i and j, Eq. (4.28) distinguishes between two partitions that differ only in that the elements of the ith and jth subsets are interchanged. For example, if $A = \{1, 2, 3, 4, 5\}$, then the partition

$$A_1 = \{2\} \quad A_2 = \{1, 4\} \quad A_3 = \{3, 5\} \tag{4.29}$$

is considered to be different from

$$A_1 = \{2\} \quad A_2 = \{3, 5\} \quad A_3 = \{1, 4\} \tag{4.30}$$

As an example, consider that we are going to split a group of 12 children into four groups of three. Equation (4.28) assumes that each group has a unique feature associated with it, so that swapping two entire groups would be considered a different outcome.

4.9.3 Applications

Example 4.9.4 The State Lottery Commission draws six balls at random from a bin of 51 numbered balls, where order does not matter in determining if someone wins. There are $\binom{51}{6}$ ways to draw these balls. Thus if you buy one lottery ticket, your chances of winning are one in about 18 million. ■

Example 4.9.5 Your club of 50 members is going to choose a president, vice-president, secretary, a banquet committee of four, and a service project committee of five. No one may serve in more than one capacity. This process can be thought of as partitioning a 50-element set into subsets of size 1, 1, 1, 4, 5, and 38. The number of ways this can be done is

$$\binom{50}{1}\binom{49}{1}\binom{48}{1}\binom{47}{4}\binom{43}{5} = \frac{50!}{1!1!1!4!5!38!} = \text{a whole bunch} \tag{4.31}$$

■

Let's put together all the counting techniques we have developed to answer some more complex questions.

Example 4.9.6 Suppose you toss a coin 10 times and observe the sequence of outcomes. From Example 4.8.14, we know that there are 2^{10} possible outcomes. We want to count the number of these outcomes that have exactly 3 heads. Imagine we have the numbers in $\mathbb{N}_{10} = \{1, \ldots, 10\}$ written on slips of paper. An outcome of the 10 coin tosses with exactly 3 heads can be uniquely specified by taking a 3-element subset of $\mathbb{N}_{10}$ to represent which of the 10 slots are heads. There are $C(10, 3) = (10 \cdot 9 \cdot 8)/(3 \cdot 2 \cdot 1) = 120$ such choices. ■

Example 4.9.7 A box of light bulbs contains 100 bulbs, 3 of which are defective. If we choose 10 bulbs from the box, we want to calculate how many ways there are to get precisely one of the defective bulbs. Such a choice of 10 bulbs consists of 9 of the 97 good bulbs chosen without regard to order, and 1 of the defective bulbs chosen from the 3 bulbs. Multiplying these, we have $C(97, 9) \times C(3, 1)$ ways of choosing the bulbs in which precisely one is defective. ■

Exercise 4.6.7 says that if A and B are disjoint, then $|A \cup B| = |A| + |B|$. Applied to counting the number of ways of performing a task, we call this theorem the *sum rule*. If counting the number of ways of performing a task must be broken up into disjoint cases, the sum rule says simply that the calculations for the different cases are added to produce the total number of ways of performing the task.

Example 4.9.8 A license plate consists of up to seven characters, taken from the 26 letters and the digits 0–9. If there is room for spaces, they are not counted in the arrangement; for example, there is no distinction between $\boxed{\text{TANGENT}}$ and $\boxed{\text{TAN GENT}}$. We count the total number of plates that can be issued. There are different cases based on the number of characters used. The number of one-character license plates is 36; the number of two-character license plates is 36^2, and so on. Thus the total number of license plates is

$$36 + 36^2 + 36^3 + \cdots + 36^7 = 36(1 + \cdots + 36^6) = 36\left(\tfrac{36^7-1}{36-1}\right) = \tfrac{36}{35}(36^7 - 1) \quad ■$$

Example 4.9.9 Your club of 30 women and 20 men is going to choose a committee of five. We calculate the number of ways to choose the committee if it must include at least one man. The number of ways to choose the committee with no restrictions is $C(50, 5)$, and the number of ways to choose the committee with no men is $C(30, 5)$. By the complement rule, the number of ways to choose the committee where there is at least one man is $C(50, 5) - C(30, 5)$. ■

Exercise 4.9.10 Pizza Peddler offers 12 different toppings and is having a special on their two-topping pizzas. How many distinct ways are there to order one of their special pizzas?

Exercise 4.9.11 You overheard that a TRUE/FALSE test of 10 questions had 7 TRUE and 3 FALSE answers. If you take the test with this knowledge, how many ways are there to fill out the answer sheet?

EXERCISE 4.9.12 How many ways are there to arrange all the letters in the word MISSISSIPPI?[20]

EXERCISE 4.9.13 Your office is 10 blocks east and 7 blocks north of your apartment, in a section of the city where the streets run uninterruptedly north-south and east-west. Every day you go to work by walking along these streets the 17 blocks, always going either east or north. How many distinct paths of this sort are there from your apartment to your office?

EXERCISE 4.9.14 When you walk home from the office (as in Exercise 4.9.13), you always like to stop at the ice cream parlor located on the corner two blocks south and five blocks west of the office. How many distinct paths back home are there if you always stop for ice cream?

EXERCISE 4.9.15 Your club consists of 50 members, 22 men and 28 women. A committee of 5 is to be chosen, which must contain 2 men and 3 women. How many distinct committees can be chosen?

EXERCISE 4.9.16 A jury of 12 and 2 alternates is chosen from 30 people summoned for jury duty. How many ways can this be done?

EXERCISE 4.9.17 Your club of 30 women and 20 men is going to choose a committee of five, which must not consist of five people of the same gender. How many ways can it be done?

EXERCISE 4.9.18 The State Lottery Commission (Example 4.9.4) is considering including the number 52 in its drawing, so that the game is played by drawing six numbers from $\mathbb{N}_{52}$. By what percentage is the number of possible outcomes increased?

EXERCISE 4.9.19 In a room with eight people, everyone shakes hands with everyone else exactly once. How many handshakes take place?

EXERCISE 4.9.20 Your freshman seminar has 12 people in it, and your professor decides to have several dinner parties for 4 of you at a time. If she wants to have a party for every possible group of 4 students from the class, how many parties will she have?

EXERCISE 4.9.21 You have a penny, nickel, dime, quarter, and half dollar, and you are going to distribute them among your 12 nieces and nephews, making sure that no child gets more than one coin. In how many ways can this be done?

[20] Generalize Example 4.9.6.

Exercise 4.9.22 You have 5 quarters, and you are going to distribute them among your 12 nieces and nephews, making sure that no child gets more than one coin. In how many ways can this be done?

Exercise 4.9.23 In a box of 50 light bulbs, 5 are known to be defective. You take a sample of 3 bulbs from the box. How many different samples contain exactly one defective bulb?

Exercise 4.9.24 Andrew, Bethany, Chigger, and Delusia are going to play spades, where all 52 cards are dealt out, 13 to each player. How many ways can the cards be dealt?

Exercise 4.9.25 You have invited five friends over for dinner, none of whom gave a firm commitment to come. How many different ways can attendance at the dinner turn out?

Exercise 4.9.26 Your first-year seminar class of 16 students is going to break up into four discussion groups of four students each. The groups are distinguishable, because each group will discuss a different question. How many ways can this be done?

Exercise 4.9.27 In the previous exercise, suppose the discussion groups are now indistinguishable because each group will discuss the same question. How many ways can this be done?

Exercise 4.9.28 Suppose you draw 5 cards from a standard deck of 52 cards. How many such draws have:

(a) At least one ace?

(b) No face cards?

(c) At least one face card?

(d) All cards of the same suit?

Exercise 4.9.29 A four-letter "word" is to be formed from the letters in the set {A, B, C, D, E, F, G}. Repetition of letters is allowed. How many such words contain:

(a) The letter E?

(b) A vowel?

Exercise 4.9.30 Rework the previous exercise under the assumption that repetition of letters is not allowed.

Exercise 4.9.31 The U.S. Senate consists of 100 senators, with 2 from each state. In how many ways may a committee of 5 senators be formed so that no 2 senators are from the same state?

Exercise 4.9.32 You have 12 nieces and nephews, 3 of whom are kids of your pyromaniac brother Torch. If you distribute 5 quarters among the 12 kids so that no child receives more than one quarter, how many such distributions have at least one of Torch's kids receiving a quarter?

Exercise 4.9.33 You have 4 novels, 6 comic books, and 9 textbooks on your bookshelf. In how many different ways can you arrange them if you want the books of each type to be grouped together?

4.10 The Binomial Theorem

The appearance of combinations in a variety of settings reveals some nice ties between seemingly unrelated mathematical ideas. First, there is a nice way to visualize Exercise 4.9.2 in *Pascal's triangle* (Fig. 4.8). To construct Pascal's triangle, we first extend the definition of $C(n, k)$ to allow for $k < 0$ and $k > n$. In both of these cases we define $C(n, k) = 0$.

Exercise 4.10.1 Verify that Eq. (4.26) holds for $k \leq 0$ and $k \geq n + 1$.

To construct Pascal's triangle, we begin with a row that corresponds to combinations of the form $C(0, k)$, which we will therefore call row zero. Because $C(0, 0) = 1$ and $C(0, k) = 0$ for all nonzero values of k, we construct row 0 to have a single 1 entry, with 0s elsewhere.

The next row corresponds to $n = 1$, and so on, and each entry in a row is determined by adding the two entries diagonally above it in the previous row.

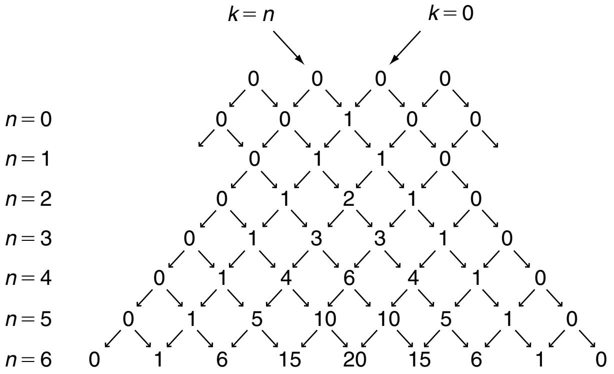

Figure 4.8 Pascal's triangle.

By Eq. (4.26), it follows that the entries in row n are $C(n, k)$, where the diagonal column of 1s headed downward left and downward right correspond to $k = 0$ and $k = n$, respectively.

Here's an important place where the entries in Pascal's triangle appear. Suppose you want to expand the expression $(a + b)^n$ for some nonnegative integer n. Rather than multiply out $(a + b)(a + b) \cdots (a + b)$ using the extended distributive property, there is a much easier way to see what the terms are, and it involves $C(n, k)$ in the coefficients.

Theorem 4.10.2 (Binomial Theorem). Let a and b be nonzero real numbers, and let n be a nonnegative integer. Then

$$\begin{aligned}(a+b)^n &= \binom{n}{0}a^n b^0 + \binom{n}{1}a^{n-1}b^1 + \binom{n}{2}a^2b^2 + \cdots + \binom{n}{n}a^0b^n \\ &= \sum_{k=0}^{n}\binom{n}{k}a^{n-k}b^k = \sum_{k=0}^{n}\binom{n}{n-k}a^{n-k}b^k\end{aligned} \tag{4.32}$$

The only reason we do not allow a or b to be zero in Theorem 4.10.2 is that 0^0 is not defined, so Eq. (4.32) would produce some undefined terms. If either a or b is zero, the expansion of $(a + b)^n$ is not particularly interesting. Thus we omit it. Using the entries from Pascal's triangle, we have

$$\begin{aligned}(a+b)^0 &= 1a^0b^0 = 1 \\ (a+b)^1 &= 1a^1b^0 + 1a^0b^1 = a + b \\ (a+b)^2 &= a^2 + 2ab + b^2 \\ (a+b)^3 &= a^3 + 3a^2b + 3ab^2 + b^3, \text{ etc.}\end{aligned} \tag{4.33}$$

We can argue the binomial theorem in several ways using combinatorial style arguments that appeal to some of our previous counting techniques. One way is to note that applying the extended distributive property to $(a + b)(a + b) \cdots (a + b)$ is tantamount to creating a whole bunch of terms of a form like *aaabbaab*, where from each $(a + b)$ factor we select either a or b. Since each factor allows for two possible choices, there are 2^n terms generated, and every one is of the form $a^{n-k}b^k$ for some $0 \le k \le n$. To determine the coefficient of $a^{n-k}b^k$, we must determine how many times the term $a^{n-k}b^k$ appears in all the distribution of the multiplication. If k of the factors supply us with b and the remaining factors provide us with a, then the number of times $a^{n-k}b^k$ appears in the grand sum is the same as the number of ways of choosing k of the n terms to provide us with b (or $n - k$ of the terms to provide us with a). That is, the term $a^{n-k}b^k$ is produced $C(n, k)$ times in the distribution.

EXERCISE 4.10.3 Determine the following.

(a) The coefficient of a^2b^5 in $(a+b)^7$

(b) The coefficient of a^2b^5 in $(a-b)^7$

(c) The coefficient of x^4 in $(x+2)^6$

(d) The coefficient of x^3y^2 in $(3x-y)^5$

We're also going to provide an algebraic proof of Theorem 4.10.2. The technique we will use in making the inductive step contains some manipulation of summations that can really come in handy in some of your other coursework, especially, for example, differential equations and complex analysis, where you must resort to finding what is called a *series solution* to a problem. Watch how we pull off some first and last terms from the summations and then realign the terms by changing the counter variable.

Proof of Theorem 4.10.2. Let a and b be nonzero real numbers.

(J1) For $n=0$, $(a+b)^0 = 1 = C(0,0)a^0b^0$, so Eq. (4.32) is true for $n=0$.

(J2) Suppose $n \geq 0$ and that Eq. (4.32) is true for n. Then

$$\begin{aligned}(a+b)^{n+1} &= (a+b)^n(a+b) = \left[\sum_{i=0}^{n} C(n,i)a^{n-i}b^i\right](a+b)\\ &= \sum_{i=0}^{n} C(n,i)a^{n+1-i}b^i + \sum_{i=0}^{n} C(n,i)a^{n-i}b^{i+1} \qquad (4.34)\\ &= \sum_{i=0}^{n} C(n,i)a^{(n+1)-i}b^i + \sum_{i=0}^{n} C(n,i)a^{(n+1)-(i+1)}b^{i+1}\end{aligned}$$

To align the terms in the two summations, let $j = i+1$ for the second summation (that is, $i = j-1$), and note that $0 \leq i \leq n$ is the same as $1 \leq j \leq n+1$. Making this substitution first, then pulling off a few individual terms, we can continue Eq. (4.34) further in several steps.

$$\begin{aligned}&= \sum_{i=0}^{n} C(n,i)a^{(n+1)-i}b^i + \sum_{j=1}^{n+1} C(n,j-1)a^{(n+1)-j}b^j\\ &= \sum_{i=1}^{n} C(n,i)a^{(n+1)-i}b^i + \sum_{j=1}^{n} C(n,j-1)a^{(n+1)-j}b^j \qquad (4.35)\\ &\quad + C(n,0)a^{n+1-0}b^0 + C(n,n)a^0b^{n+1}\end{aligned}$$

Now $C(n, 0) = C(n+1, 0)$ and $C(n, n) = C(n+1, n+1)$. Also, we may let $k = i = j$ and combine the two summations and apply Exercise 4.9.2 to have

$$\begin{aligned}
&= \sum_{k=1}^{n} [C(n, k) + C(n, k-1)] a^{(n+1)-k} b^k \\
&\qquad + C(n+1, 0) a^{n+1-0} b^0 + C(n+1, n+1) a^0 b^{n+1} \\
&= \sum_{k=1}^{n} C(n+1, k) a^{(n+1)-k} b^k \qquad (4.36) \\
&\qquad + C(n+1, 0) a^{n+1-0} b^0 + C(n+1, n+1) a^0 b^{n+1} \\
&= \sum_{k=0}^{n+1} C(n+1, k) a^{(n+1)-k} b^k
\end{aligned}$$

which verifies that Eq. (4.32) holds for $n + 1$. □

The algebraic proof of Theorem 4.10.2 means that we have provided an algebraic proof of the following.

Corollary 4.10.4 For all nonnegative integers n,

$$\sum_{k=0}^{n} \binom{n}{k} = 2^n \qquad (4.37)$$

Proof. Let $a = b = 1$ in Eq. (4.32). □

Suppose A is a finite set with cardinality n. Corollary 4.10.4 implies that A has 2^n subsets, for the left-hand side of Eq. (4.37) is the sum of the number of subsets of an n-element set with all possible numbers of elements. Another way to see the same result is to create a one-to-one correspondence between the family of all subsets of A and the collection of the 2^n functions from A to $\{0, 1\}$. If A_1 is a given subset of A, we may define

$$f(a) = \begin{cases} 1, & \text{if } a \in A_1 \\ 0, & \text{if } a \notin A_1 \end{cases} \qquad (4.38)$$

Every distinct subset of A generates a distinct function, and by Theorem 4.8.8, there are 2^n functions from A to $\{0, 1\}$.

Exercise 4.10.5 Use your imagination to generalize the binomial theorem, so that you can determine the following.

(a) The coefficient of $a^2b^2c^2$ in $(a + b + c)^6$

(b) The coefficient of a^3b^3 in $(a + b + c)^6$

(c) The coefficient of ab^3c^2 in $(a - b + c)^6$

(d) The coefficient of a^3b^3 in $(a - b - c)^6$

(e) The coefficient of a^3b^2 in $(a + b + 5)^7$

(f) The coefficient of ab^4 in $(a + 2b + 3)^8$

(g) The coefficient of $a^2b^2c^3$ in $(2a + b - 3c + 5)^{10}$

This page intentionally left blank

PART II

Basic Principles of Analysis

This page intentionally left blank

The Real Numbers

Let's look at some defining characteristics of the area of mathematics we call analysis. Given a set S, its elements might be endowed with a measure of size. The size, or *norm*, of an element x is typically denoted $|x|$, like absolute value on the real numbers, or perhaps $\|x\|$. The norm of an element will always be a nonnegative real number, so that a norm is really just a function from S into the real numbers that has three defining characteristics. You have already run into these in Section 2.3. They are properties N1–N2 on page 54 and property N3 on page 55.

Measuring the sizes of elements is only one type of structure that can be placed on a set that puts it squarely in the field of analysis, but some notion of measure with nonnegative real numbers is characteristic of structures in analysis. For example, some structures do not have a norm but have a way of measuring some idea of distance between elements. Such a measure of distance is called a *metric*. Whether S is endowed with a norm or a metric, the measure it represents is inextricably tied to the real numbers. So the set of real numbers is at the heart of analysis, and one could argue that no analysis into any structure except the real numbers should be undertaken until one understands the axioms and fundamental results of the theory of real numbers.

Be that as it may, a norm or metric on S, if one exists, lends itself to much fruitful study: Sequences and their convergence, continuity, and calculus are but a few. In this chapter, we address more advanced properties of the real numbers that arise out of the assumptions from Chapter 0, in particular A20. All these properties are important not only because they apply to the real numbers, but also because they are typical of properties of many other structures in analysis. In this chapter and the one to follow, all sets are assumed to be subsets of the real numbers unless specified otherwise.

5.1 The Least Upper Bound Axiom

The least upper bound (LUB) axiom is a standard axiom of the real numbers, endowing it with some of its familiar features. For example, the way we visualize the real numbers as a smooth number line with no holes is due to the LUB axiom. Two other characteristics of the real numbers are logically equivalent to the LUB axiom. One is called the *Nested Interval Property* (NIP) and has to do with what you get when you intersect a whole bunch of real number intervals that have certain properties. The other is called *completeness* and has to do with the way certain sequences behave when the terms get close to each other. It is possible to assume any one of the LUB axiom, the NIP, or completeness and derive the other two as theorems. Assuming the LUB property as an axiom is probably most common, so we choose that route. In this chapter, we will prove the NIP from the LUB axiom and completeness from the NIP, and look into the converses of these theorems to show that all three are logically equivalent. In this section we explore the LUB property in some depth to get a feel for what it means, and we derive some immediately important implications. First, let's present a definition.

Definition 5.1.1 A set A is said to be *bounded from above* if there exists M_1 such that $a \le M_1$ for all $a \in A$. A is said to be be *bounded from below* if there exists M_2 such that $a \ge M_2$ for all $a \in A$. A is said to be *bounded* if there exists $M > 0$ such that $|a| \le M$ for all $a \in A$.

The first exercise appears to be a very obvious result, but it must be demonstrated in terms of Definition 5.1.1. Proving the ⇒ direction is easier. For if A is bounded, then the guaranteed $M > 0$ such that $|a| \le M$ for all $a \in A$ should clearly suggest values of M_1 and M_2. However, in proving the ⇐ direction, you must use the existence of M_1 and M_2 to create a single, positive value of M such that $-M \le a \le M$ for all $a \in A$.

EXERCISE 5.1.2 A set of real numbers is bounded if and only if it is bounded from above and below.[1]

5.1.1 Least Upper Bounds

Suppose A is non-empty and bounded from above. Among all possible upper bounds for A, we would call L a *least upper bound* for A if L is an upper bound for A with the additional property that $L \le N$ for every upper bound N. That is, among all upper bounds for A, none is any smaller than L. The LUB axiom says that every non-empty set of real numbers that is bounded from above has an LUB that is a real number. We state it again here for the sake of reference.

(A20) **Least upper bound property:** If A is a non-empty subset of real numbers that is bounded from above, then there exists a real number L with the following properties:

(L1) For every $a \in A$, we have that $a \le L$, and

(L2) If N is any upper bound for A, it must be that $N \ge L$.

[1] To prove ⇐, let $M = |M_1| + |M_2| + 1$. Use Exercise 2.3.7.

EXERCISE 5.1.3 The least upper bound of a set of real numbers is unique.

In this chapter we will make extensive use of interval notation to represent sets of real numbers.

$$(a, b) = \{x \in \mathbb{R} : a < x < b\} \tag{5.1}$$

$$[a, b] = \{x \in \mathbb{R} : a \leq x \leq b\} \tag{5.2}$$

$$(a, +\infty) = \{x \in \mathbb{R} : x > a\} \tag{5.3}$$

$$(-\infty, a] = \{x \in \mathbb{R} : x \leq a\} \tag{5.4}$$

Intervals of the form (5.1) and (5.3) are called *open* intervals, and those of the form (5.2) and (5.4) are called *closed* intervals. The motivation for these terms will become clear in Section 5.3 where we discuss open and closed sets in general.

Example 5.1.4 Suppose $a < b$. Show that b is the LUB of the interval (a, b).

Solution If $x \in (a, b)$, then $x < b$, so that b has property L1. To show that b has property L2, let $N < b$ be any real number. We show that N is not an upper bound of (a, b). If $N < a$, then by Exercise 2.2.9,

$$N < a < (a + b)/2 < b$$

so that $(a + b)/2$ is an element of (a, b) that is greater than N. On the other hand, if $N \geq a$, then since $N < b$ also, we have

$$a \leq N < (N + b)/2 < b$$

Thus $(N + b)/2$ is an element of (a, b) that is strictly greater than N. In either case, N is not an upper bound of (a, b), and b therefore has property L2. ■

Least upper bounds can be defined in terms of an alternate pair of properties, which we will call M1–M2. In some situations, showing a particular number is the LUB of a set is more easily done by showing that it has these alternate properties. Your task (Exercise 5.1.6) will be to show that properties L1–L2 are equivalent to M1–M2. But first, let's state the equivalent form as a theorem to suggest a way of attacking its proof.

Theorem 5.1.5 Given a set A, L is the LUB for A if and only if the following two conditions hold.

(M1) For every $\epsilon > 0$ (no matter how large), $(L, L + \epsilon) \cap A$ is empty.

(M2) For every $\epsilon > 0$ (no matter how small), there exists $a \in A \cap (L - \epsilon, L]$.

The Greek letter ϵ (epsilon) has been used so much in analysis to represent an arbitrary positive real number that it has come to have a personality all its own.

Although ϵ is not generally thought to be of any particular size, it is usually present in theorems and proofs because smaller values of ϵ represent the primary obstacle to overcome in concocting the proof. Theorem 5.1.5 is a natural way to visualize the LUB property in terms of how intervals to the left and right of L intersect the set A. Every interval of the form $(L, L+\epsilon)$, no matter how large ϵ is, must not contain any elements of A. Also, every interval of the form $(L-\epsilon, L]$, no matter how small ϵ is, must contain some element of A.

Here is a suggestion about how to tackle the proof of Theorem 5.1.5. Given a set A and a real number L, the theorem says (L1 ∧ L2) ↔ (M1 ∧ M2). For the → direction, Exercises 1.2.18(f) and (i) imply that

$$\begin{aligned}(\text{L1} \wedge \text{L2}) \to (\text{M1} \wedge \text{M2}) &\Leftrightarrow [(\text{L1} \wedge \text{L2}) \to \text{M1}] \wedge [(\text{L1} \wedge \text{L2}) \to \text{M2}] \\ &\Leftrightarrow [(\text{L2} \wedge \neg\text{M1}) \to \neg\text{L1}] \wedge [(\text{L1} \wedge \neg\text{M2}) \to \neg\text{L2}]\end{aligned} \tag{5.5}$$

The statement ¬M1 → ¬L1 is at least as strong as (L2 ∧ ¬M1) → ¬L1. So to prove the → direction of Theorem 5.1.5, you will want to show L1 → M1 by contrapositive and then show (L1 ∧ L2) → M2 by showing (L1 ∧ ¬M2) → ¬L2. Then you can prove ← by showing ¬L1 → ¬M1 and ¬L2 → ¬M2.

EXERCISE 5.1.6 Prove Theorem 5.1.5.

Example 5.1.7 Show that the set {1} has LUB 1 by appealing to properties M1–M2.

Solution Pick $\epsilon > 0$. Because $(1, 1+\epsilon) \cap \{1\}$ is empty, property M1 holds. Also, because $1 \in (1-\epsilon, 1] \cap \{1\}$, property M2 holds. Thus 1 is the LUB of {1}. ■

EXERCISE 5.1.8 Show that b is the LUB of the interval (a, b) by appealing to properties M1–M2. (An identical argument will work for $[a, b]$.)

5.1.2 Greatest Lower Bounds

Now let's turn the LUB property upside down and discuss bounds from below. We do not have to make any new assumptions concerning the existence of the greatest lower bound of a set, for we can derive results from the LUB property by flipping a set upside down, so to speak.

Definition 5.1.9 A real number G is said to be a *greatest lower bound* (GLB) for a set A if G has the following properties:

(G1) For every $a \in A$, we have that $a \geq G$, and

(G2) If N is any lower bound for A, it must be that $N \leq G$.

Exercise 5.1.10 The greatest lower bound of a set is unique.

The next exercise will serve as a helpful lemma for the exercise to follow.

Exercise 5.1.11 Suppose A is a set that is bounded from below by M, and let $B = \{-a : a \in A\}$. Then B is bounded from above by $-M$.

Exercise 5.1.12 If A is a non-empty set of real numbers that is bounded from below, then A has a GLB.[2]

Exercise 5.1.13 State and prove a theorem analogous to Theorem 5.1.5 for the GLB of a set.

Exercise 5.1.14 Let A be a subset of the real numbers, and define the set $B = \{-a : a \in A\}$. Then A has LUB L if and only if B has GLB $-L$.[3]

5.2 The Archimedean Property

The LUB property implies that the positive integers are unbounded from above, which is the claim of the next exercise. To prove it, suppose the result is false, claim the existence of a LUB, and then exploit properties M1–M2 for a strategically chosen value of ϵ to produce a contradiction.

Exercise 5.2.1 For all positive real numbers x, there exists a positive integer n such that $n > x$.

Exercise 5.2.1 has a logically equivalent form called the *Archimedean property*. We have already pointed out briefly that the ancient Greeks were very sophisticated mathematically. One idea they used involved what we would call an *infinitesimal* number. Different from zero, an infinitesimal was considered to be smaller than every positive number. One salient property of an infinitesimal is that you can add it to itself any finite number of times and still have an infinitesimal. The next result says that there are no infinitesimals in the real numbers. It is named after Archimedes, who argued against the use of infinitesimals.

Exercise 5.2.2 [Archimedean Property] For all positive real numbers ϵ and r, there exists a positive integer n such that $n\epsilon > r$.[4]

Thus no matter how small a positive number ϵ might be, and no matter how large a positive number r is, you can add ϵ to itself some n number of times

[2] If B from Exercise 5.1.11 has LUB L, you can show that $-L$ is the GLB for A.

[3] You might be able to appeal to your proof of Exercise 5.1.12.

[4] Apply Exercise 5.2.1 to $x = r/\epsilon$.

to produce a sum that exceeds r, so that ϵ is therefore not an infinitesimal. The Archimedean property often proves to be useful in a slightly different form, using the specific case $r = 1$. We state it as a corollary to Exercise 5.2.2.

Corollary 5.2.3 For all $\epsilon > 0$, there exists a positive integer n such that $1/n < \epsilon$.

This alternate form of the Archimedean property says that numbers of the form $1/n$ can do a pretty mean limbo dance in the real numbers. No matter how low you set the ϵ-bar, the reciprocal of some positive integer can wiggle under it.

Exercise 5.2.4 Between any two real numbers $a < b$, there exists a *nonzero* rational number.[5,6,7,8]

Exercise 5.2.5 Between any two real numbers $a < b$, there exists an irrational number.[9]

5.2.1 Maximum and Minimum of Finite Sets

The intervals (a, b) and $[a, b]$ illustrate that the LUB and GLB of a set might or might not actually be in the set. If a set is finite, then the LUB and GLB will be elements of the set, and we will call them the *maximum* and *minimum*, respectively. The Archimedean property plays an important role in the proofs of these facts.

Exercise 5.2.6 If $S = \{a_k\}_{k=1}^{p}$ is a finite set of real numbers, then S is bounded.[10]

Which of the following statements is stronger?

$$\text{For every disease, there exists a medicine that cures that disease.} \tag{5.6}$$

or

$$\text{There exists a medicine that cures every disease.} \tag{5.7}$$

Statement (5.7) says that there exists a panacea, so that clearly (5.6) is true. But just because (5.6) is true, it might be because there are many different medications, each of which applies only to one disease. So a panacea does not necessarily exist, and (5.7) is not necessarily true. Thus (5.7) is stronger than (5.6).

[5] First start by assuming $a > 0$. Worry about the other cases later.

[6] Apply Corollary 5.2.3 to $\epsilon = b - a$.

[7] If $1/n < b - a$, you can apply Exercise 5.2.2 to a and $1/n$.

[8] Use the WOP for the right m such that $m/n > a$.

[9] Apply the result of Exercise 5.2.4 to $a/\sqrt{2} < b/\sqrt{2}$.

[10] Let $M = 1 + \sum_{k=1}^{p} |a_k|$.

Now suppose we have a finite set $\{a_k\}_{k=1}^{p}$ in which each a_k is positive. Then we may apply the Archimedean property to each a_k separately to claim the existence of positive integers $\{n_1, n_2, \ldots, n_p\}$, where $1/n_k < a_k$ for every k. This is like saying every disease has a cure, because it says every a_k has its own n_k. But in the next exercise, we claim that there is a single positive integer n that satisfies $1/n < a_k$ for all k. To find this panacea n, you will need to build it in terms of all the n_k.

EXERCISE 5.2.7 Let $S = \{a_k\}_{k=1}^{p}$ be a finite set of *positive* real numbers. Then there exists a positive integer n such that $1/n < a_k$ for all $1 \le k \le p$.[11]

The Archimedean property says that numbers of the form $1/n$ are clustered around zero arbitrarily closely. If you zoom in to any interval $(-\epsilon, \epsilon)$, no matter how small ϵ might be, there will be some $1/n$ in this interval. You cannot zoom in close enough to find an interval around zero that is devoid of numbers of the form $1/n$. In general, if it is possible to find some $\epsilon > 0$ such that the interval $(L - \epsilon, L + \epsilon)$ contains no elements of a given set, we say the set can be *bounded away from* L. Exercise 5.2.7 says that a finite set of positive numbers can be bounded away from zero. We state this fact as a corollary.

Corollary 5.2.8 A finite set of positive real numbers can be bounded away from zero. That is, if $S = \{a_k\}_{k=1}^{p}$ is a finite set of positive real numbers, then there exists $M > 0$ such that $a_k > M$ for all $1 \le k \le p$.

If S is a finite set, then by Exercise 5.2.6, S is bounded. Thus by Exercise 5.1.2, it is bounded from above and below. By the LUB and GLB properties, S has a LUB and GLB. In fact:

Theorem 5.2.9 A finite set of real numbers contains its LUB and GLB.

EXERCISE 5.2.10 Prove that the set in Theorem 5.2.9 contains its LUB. (The argument that it contains its GLB would be similar.)[12]

Definition 5.2.11 Let S be a finite set of real numbers. Then the LUB and GLB of S are called the *maximum* and *minimum* values of S, respectively, and are denoted $\max(S)$ and $\min(S)$.

The maximum and minimum of a finite set will prove to be important in later sections. Here are two simple illustrations. First, suppose anyone at least 16 years of age can drive, anyone at least 18 can vote, and anyone at least 65 can collect Social Security. How can you guarantee that a chosen person can drive, vote, and collect Social Security? Certainly you would choose a person whose age is at least

[11] Let $n = \sum_{k=1}^{p} n_k$.

[12] If L is the LUB of S and $L \notin S$, then you ought to be able to use the set $T = \{L - a_k\}_{k=1}^{p}$ and Corollary 5.2.8 to contradict the assumption that L is the LUB of S.

$\max\{16, 18, 65\}$. For if we write $M = \max\{16, 18, 65\}$, then a person of age M can drive, vote, and collect Social Security because M satisfies all three inequalities $M \geq 16, M \geq 18$, and $M \geq 65$.

As another example, suppose every real number within ϵ_1 distance of zero is in set A, and every real number within ϵ_2 of zero is in set B. That is, if $|x| < \epsilon_1$, then $x \in A$, and if $|x| < \epsilon_2$, then $x \in B$. If we let $\epsilon = \min\{\epsilon_1, \epsilon_2\}$, we can be sure that every real number within ϵ of zero will be in $A \cap B$. For if $|x| < \epsilon$, then both $|x| < \epsilon_1$ and $|x| < \epsilon_2$ are true, so that $x \in A$ and $x \in B$.

The terms $\max(S)$ and $\min(S)$ are often used when S is not finite in the event that S contains its LUB and GLB.

5.3 Open and Closed Sets

Research in the area of *topology* flourished more through the middle half of the twentieth century than in any other period. In these next few sections, we introduce some of the basic terminology that is commonly used in topology, but we apply them only to the real numbers.

Definition 5.3.1 Given a real number a and $\epsilon > 0$, we define the *ϵ-neighborhood* of a as

$$N_\epsilon(a) = (a - \epsilon, a + \epsilon) = \{x \in \mathbb{R} : a - \epsilon < x < a + \epsilon\} = \{x \in \mathbb{R} : |x - a| < \epsilon\} \tag{5.8}$$

and we say that the neighborhood has *radius* ϵ.

Notice if $0 < \epsilon_1 \leq \epsilon_2$, then $N_{\epsilon_1}(a) \subseteq N_{\epsilon_2}(a)$. For if $x \in N_{\epsilon_1}(a)$, then

$$a - \epsilon_2 \leq a - \epsilon_1 < x < a + \epsilon_1 \leq a + \epsilon_2 \tag{5.9}$$

so that $x \in N_{\epsilon_2}(a)$.

Definition 5.3.2 A set A is called *open* provided for every $a \in A$, there exists $\epsilon > 0$ such that $N_\epsilon(a) \subseteq A$.

A set is open if every element is in some interval contained entirely within the set; that is, every point is, in a sense, insulated from the outside of the set by a neighborhood that is contained completely within the set.

EXERCISE 5.3.3 Show that the intervals $(a, +\infty)$ and $(-\infty, b)$ are open.

What does it mean for a set A not to be open?

$$\begin{aligned} \neg(A \text{ is open}) &\Leftrightarrow \neg(\forall a \in A)(\exists \epsilon > 0)(N_\epsilon(A) \subseteq A) \\ &\Leftrightarrow (\exists a \in A)(\forall \epsilon > 0)(N_\epsilon \not\subseteq A) \qquad (5.10) \\ &\Leftrightarrow (\exists a \in A)(\forall \epsilon > 0)(\exists x \in N_\epsilon(a))(x \notin A) \end{aligned}$$

Thus A is not open provided there is at least one point in A, for which every ϵ-neighborhood contains some point not in A. Some point of A has points of A^C in every ϵ-neighborhood, no matter how small ϵ might be.

EXERCISE 5.3.4 Show that the singleton set $\{a\}$ is not open.

EXERCISE 5.3.5 Show that the set $(a, b]$ is not open.

Two other sets besides $(a, +\infty)$ and $(-\infty, b)$ immediately present themselves as open sets. The empty set is open simply because otherwise there would have to exist some x in the empty set with some property, and that is a contradiction. The real numbers are an open set because for any real number a, we may let $\epsilon = 1$ and have $N_\epsilon(a) \subseteq \mathbb{R}$.

Let's address unions and intersections of open sets. If A and B are both open, do you suspect $A \cup B$ is open? What about $A \cap B$? In both cases the answer is yes, but when an arbitrary family of open sets is combined by union or intersection, the answer can change.

EXERCISE 5.3.6 Let $\mathcal{F} = \{A\}$ be a family of open sets, and let $\{B_k\}_{k=1}^n$ be a finite family of open sets. Then

(a) $\bigcup_{\mathcal{F}} A$ is open.

(b) $\bigcap_{k=1}^n B_k$ is open.

By Exercise 5.3.6, it follows that an interval of the form (a, b) is open, for $(a, b) = (-\infty, b) \cap (a, +\infty)$. The intersection of an infinite family of open sets certainly could be open, but it might not be.

EXERCISE 5.3.7 Show that the intersection of infinitely many open sets might not be open by showing $\bigcap_{n=1}^\infty (-1/n, 1/n) = \{0\}$.[13]

If we were discussing doors or stores or minds, to say that one is closed would mean that it is not open. However, *to say a set is closed does not mean it is not open.*

Definition 5.3.8 A set A is said to be *closed* provided A^C is open.

EXERCISE 5.3.9 The singleton set $\{a\}$ is closed.

EXERCISE 5.3.10 If $a < b$, then the interval $[a, b]$ is closed.

[13] Show that zero is an element of every interval $(-1/n, 1/n)$ and that any nonzero real number fails to be in some $(-1/n, 1/n)$.

EXERCISE 5.3.11 Let $\mathcal{F} = \{A\}$ be a family of closed sets, and let $\{B_k\}_{k=1}^n$ be a finite family of closed sets. Then

(a) $\bigcap_{\mathcal{F}} A$ is closed.[14]

(b) $\bigcup_{k=1}^n B_k$ is closed.

By Exercises 5.3.9 and 5.3.11, it follows that any finite set is closed because it is a finite union of singleton sets. The integers are a good example of an infinite set that is closed, because $\mathbb{Z}^C = \cup_{n\in\mathbb{Z}}(n, n+1)$.

EXERCISE 5.3.12 Give an example of an infinite family $\mathcal{F}$ of closed sets such that $\cup_{\mathcal{F}} A$ is not closed. Prove.[15]

EXERCISE 5.3.13 Show that $\{1/n\}_{n=1}^{\infty}$ is not closed by showing its complement is not open.

Lots of sets are neither open nor closed. For example, $[a, b)$ is sometimes called a *half-open* or *half-closed* interval. Believe it or not, some sets are both open and closed. The empty set and the real numbers are both open and closed because they are both open and they are complementary. Theorem 5.3.14 uses the LUB axiom to show the important fact that these are the only two sets of real numbers that are both open and closed. The proof is a little intricate, so we provide it here. It is important for you to work your way through the details because it illustrates an important type of proof where we show that certain examples of something are the only ones that exist. Keep a pencil and paper handy to sketch some number lines and help you visualize the details.

Theorem 5.3.14 The empty set and the real numbers are the only subsets of the real numbers that are both open and closed.

Proof. We prove by contradiction. Suppose there exists a non-empty proper subset A of the real numbers that is both open and closed. Then A and A^C are both open and non-empty. Pick any $a_1 \in A$ and $a_2 \in A^C$. Without any loss of generality, we may assume $a_1 < a_2$. Let $E = \{\epsilon > 0 : N_\epsilon(a_1) \subseteq A\}$; that is, E is the set of all $\epsilon > 0$ such that the ϵ-neighborhood of a_1 is contained entirely within A. Since A is open, E is non-empty. Because $a_2 \in A^C$, E is bounded from above by $a_2 - a_1$. Thus we may apply the LUB property to E to conclude the existence of some ϵ_0, the LUB of E. Thus $N_{\epsilon_0}(a_1)$ is the largest neighborhood of a_1 that is contained entirely within A.

In order to arrive at a contradiction, we look at the end points of the interval $(a_1 - \epsilon_0, a_1 + \epsilon_0)$. First we ask if it is possible that $a_1 + \epsilon_0 \in A^C$. If so, then the fact

[14] See Exercise 3.3.8.

[15] You shouldn't have to look too far.

that A^C is open means there exists $\epsilon_1 > 0$ such that $N_{\epsilon_1}(a_1 + \epsilon_0) \subseteq A^C$. Because $a_1 \in A$, we know that $\epsilon_1 < \epsilon_0$. But then $a_1 + \epsilon_0 - \epsilon_1/2 \in A^C$. This contradicts the fact that $N_{\epsilon_0}(a_1) \subseteq A$. Thus $a_1 + \epsilon_0 \in A$. By an identical argument applied to $a_1 - \epsilon_0$, we have that $a_1 - \epsilon_0 \in A$, so that $[a_1 - \epsilon_0, a_1 + \epsilon_0] \subseteq A$.

Now let's look again at the end points of $[a_1 - \epsilon_0, a_1 + \epsilon_0]$. Since A is open and $a_1 - \epsilon_0, a_1 + \epsilon_0 \in A$, there exist $\epsilon_2, \epsilon_3 > 0$ such that $N_{\epsilon_2}(a_1 - \epsilon_0) \subseteq A$ and $N_{\epsilon_3}(a_1 + \epsilon_0) \subseteq A$. Letting $\epsilon_4 = \min\{\epsilon_2, \epsilon_3\}$, we have that $N_{\epsilon_0+\epsilon_4}(a_1) \subseteq A$. This contradicts the fact that ϵ_0 is the LUB of E, for we have shown that $\epsilon_0 + \epsilon_4 \in E$. Thus it is impossible that A and A^C are both open and non-empty. □

Theorem 5.3.14 is important both in its own right and because it motivates the idea of *connectedness* of sets in topology. By Theorem 5.3.14, if A is a non-empty, proper, open subset of the real numbers, then A^C is a non-empty, proper subset of the real numbers that is *not* open. Thus Theorem 5.3.14 implies the following.

Corollary 5.3.15 The real numbers cannot be written as the disjoint union of two non-empty open sets.

Another way to say this is that if A and B are open and disjoint, and if $A \dot{\cup} B = \mathbb{R}$, then either $A = \emptyset$ or $B = \emptyset$. This property gives us the imagery of the real number line as one connected piece, and as the proof reveals, follows from the LUB axiom. In the more abstract setting of topology, a set is said to be connected *by definition* if it cannot be written as the disjoint union of two non-empty, open sets.

5.4 Interior, Exterior, Boundary, and Cluster Points

Suppose A is a subset of the real numbers and x is any real number. We can characterize the relationship of x to A in two distinct ways. The first is by placing x in precisely one of the interior, exterior, or boundary of A. The other is by determining whether x is a cluster point of A, which has to do with whether elements of A are packed tightly around x.

5.4.1 Interior, Exterior, and Boundary

Given a set A and any real number x, precisely one of the following must be true.

1. There exists $\epsilon > 0$ such that $N_\epsilon(x) \subseteq A$, in which case we say that x is in the *interior* of A, written $x \in \text{Int}(A)$.
2. There exists $\epsilon > 0$ such that $N_\epsilon(x) \subseteq A^C$, in which case we say that x is in the *exterior* of A, written $x \in \text{Ext}(A)$.
3. For all $\epsilon > 0$, there exists $a_1, a_2 \in N_\epsilon(x)$ such that $a_1 \in A$ and $a_2 \in A^C$, in which case we say that x is on the *boundary* of A, written $x \in \text{Bdy}(A)$.

Here are some immediate observations.

1. $\text{Int}(A) \subseteq A$. For if $x \in \text{Int}(A)$, then there is an ϵ-neighborhood of x contained entirely within A, and x is in this neighborhood. Thus $x \in A$.
2. $\text{Ext}(A) \subseteq A^C$ by similar reasoning.
3. If $x \in \text{Bdy}(A)$, then x might or might not be in A. However, if $x \in A$, then every neighborhood of x contains elements of A^C. Similarly, if $x \notin A$, then every neighborhood of x contains elements of A.
4. $\text{Int}(A) = \text{Ext}(A^C)$.
5. $\text{Ext}(A) = \text{Int}(A^C)$.
6. $\text{Bdy}(A) = \text{Bdy}(A^C)$.

EXERCISE 5.4.1 For each of the following sets, determine the interior, exterior, and boundary.

(a) $\{1\}$

(b) $(0, 1]$

(c) $\{1/n\}_{n=1}^{\infty}$

(d) $\emptyset$

(e) $\mathbb{Z}$

(f) $\mathbb{R}$

(g) $\mathbb{Q}$

The proof of the next exercise is quick. It is useful because it gives us an equivalent way to think about open sets.

EXERCISE 5.4.2 A set A is open if and only if $\text{Int}(A) = A$.

EXERCISE 5.4.3 If L is the LUB of a set A, then $L \in \text{Bdy}(A)$.

5.4.2 Cluster Points

Once again, let A be a subset of the real numbers, and let x be any real number. It might be that elements of A are packed densely around x.

Definition 5.4.4 A real number x is said to be a *cluster point* of A, provided every ϵ-neighborhood of x contains a point in A other than x itself. That is, for all $\epsilon > 0$, there exists $a \in A \cap N_\epsilon(x)$, where $a \neq x$.

A cluster point of A might or might not be an element of A.

EXERCISE 5.4.5 Show that zero is a cluster point of the following sets.

(a) $\{1/n\}_{n=1}^{\infty}$

(b) $[0, 1]$

The definition of cluster point motivates a new term for convenience.

Definition 5.4.6 For a real number x and $\epsilon > 0$, the set $N_\epsilon(x) - \{x\}$ is called the *deleted ϵ-neighborhood* of x and is denoted $DN_\epsilon(x)$. Another way to write $DN_\epsilon(x)$ is $\{y \in \mathbb{R} : 0 < |y - x| < \epsilon\}$.

Thus we may say that x is a cluster point of A if every deleted ϵ-neighborhood of x contains an element of A; that is, for every $\epsilon > 0$, $A \cap DN_\epsilon(x)$ is nonempty. Cluster points have other common names, for example, *limit points* or *accumulation points*.

EXERCISE 5.4.7 Suppose L is the LUB of a set A and $L \notin A$. Then L is a cluster point of A.

EXERCISE 5.4.8 What does it mean for x not to be a cluster point of A?

If $x \in A$, but x is not a cluster point of A, we say that x is an *isolated* point of A. Notice in the proof of the next theorem how we must replace an arbitrarily chosen $\epsilon > 0$ with $\epsilon_2 > 0$, which might be smaller. The desired point that we find in the ϵ_2-neighborhood will therefore also be in the ϵ-neighborhood.

Theorem 5.4.9 If A is an open set, then every element of A is a cluster point.

Proof. Suppose A is open and pick any $x \in A$. To show x is a cluster point of A, we must first choose $\epsilon > 0$. Since A is open, there exists $\epsilon_1 > 0$ such that $N_{\epsilon_1}(x) \subseteq A$. Let $\epsilon_2 = \min\{\epsilon, \epsilon_1\}$. Then the point $x + \epsilon_2/2 \in A \cap DN_\epsilon(x)$. □

EXERCISE 5.4.10 Show that the converse of Theorem 5.4.9 is not true.

Sometimes one statement might seem stronger than another but is actually equivalent. In the next exercise, the fact that every deleted neighborhood contains some point of A is actually strong enough to allow you to show that every deleted neighborhood contains infinitely many points of A.

EXERCISE 5.4.11 If x is a cluster point of A, then every ϵ-neighborhood of x contains infinitely many elements of A.[16]

[16] Try a proof by contrapositive.

The next exercise gives us a way to think about a closed set without having to refer to its complement.

Exercise 5.4.12 A set is closed if and only if it contains all its cluster points.

Exercise 5.4.13 Show that $\{1/n\}_{n=1}^{\infty}$ is not closed.

5.5 Closure of Sets

If a set A is not closed, we might want to close it off, so to speak, by finding the smallest closed superset of A, if there is one. First we define a term for this smallest closed superset of A. Then we address whether it exists, and if so, whether it is unique.

Definition 5.5.1 Suppose A is a set of real numbers, and suppose C is a set with the following properties.

(C1) $A \subseteq C$.

(C2) C is closed.

(C3) If D is a closed set and $A \subseteq D$, then $C \subseteq D$.

Then C is called a *closure* of A, and is denoted $\overline{A}$.

Notice how Definition 5.5.1 lays down the characteristics that the set C must have in order for it to qualify as a smallest closed superset of A. Property C1 guarantees that any set we would call $\overline{A}$ is indeed a superset of A. Property C2 guarantees that it is closed, and property C3 says that no closed superset of A will be any smaller.

Let A be a set, and consider the family of subsets of the real numbers, where each set in the family is a closed superset of A. That this family is non-empty is immediate, for $\mathbb{R}$ itself is a closed superset of A. In the next exercise, you will show that the closure of a set exists uniquely. What you will show is that the intersection of all closed supersets of A has properties C1–C3 and that any two sets that both have properties C1–C3 are actually the same.

Exercise 5.5.2 Suppose A is a set of real numbers. Then $\overline{A}$ exists uniquely and can be constructed as

$$\overline{A} = \bigcap_{\substack{S \supseteq A \\ S \text{ closed}}} S^{17} \tag{5.11}$$

[17] See Exercises 3.3.14 and 5.3.11 and Theorem 3.3.11.

The construction of $\overline{A}$ in Exercise 5.5.2 is a way of closing off a set A by intersecting down from above in a sense, through closed sets containing A, as if we were shrink wrapping A. Another way to imagine the closure of a set is in terms of its interior and boundary.

Recall that for a given set, every real number is in precisely one of the interior, exterior, or boundary of the set. Also, no element of the set is in the exterior. Exercise 5.5.4 says that the closure of a set can also be thought of as the set itself, some points of which might be interior points and some boundary points, with any missing boundary points tossed in. A hint will reveal one way to approach the proof. First, you will need the following.

EXERCISE 5.5.3 What does it mean to say $x \notin \overline{A}$?

EXERCISE 5.5.4 If A is a set of real numbers, then $\overline{A} = \text{Int}(A) \cup \text{Bdy}(A)$.[18]

EXERCISE 5.5.5 Use either Exercise 5.5.2 or Exercise 5.5.4 to determine the closure of the following sets.

(a) $\emptyset$

(b) $\mathbb{R}$

(c) (a, b)

(d) $(a, b) \cup (b, c)$

(e) $\{1/n\}_{n=1}^{\infty}$

(f) $\mathbb{Q}$

EXERCISE 5.5.6 Give an example of two disjoint sets with the same closure.

EXERCISE 5.5.7 If $A \subseteq B$, then $\overline{A} \subseteq \overline{B}$.[19]

EXERCISE 5.5.8 A is closed if and only if $\overline{A} = A$.

Corollary 5.5.9 $\overline{\overline{A}} = \overline{A}$.

Proof. Since $\overline{A}$ is closed, then by Exercise 5.5.8, $\overline{\overline{A}} = \overline{A}$. □

EXERCISE 5.5.10 One of the following statements is true, and for the other, one direction of subset inclusion is true. Prove the three true subset inclusion statements and provide a counterexample to illustrate that the fourth is false.

[18] Show $(\overline{A})^C = \text{Ext}(A)$.

[19] What is the relationship of $\overline{B}$ to A?

(a) $\overline{A \cup B} = \overline{A} \cup \overline{B}$[20]

(b) $\overline{A \cap B} = \overline{A} \cap \overline{B}$

5.6 Compactness

A special class of set deserves our attention now. Sets that are *compact* have a fundamental feature that makes them very important if they serve as the domain of certain types of functions. Before we can define compactness, we need another term first, motivated by a little story.

My greatgrandmother was known to make quilts. The standard way of making them was to collect hundreds of small pieces of cloth and then stitch them together into an interesting pattern to form a single entity—the quilt. Now suppose you are lying on the bed, and your grandmother takes a finished quilt and tosses it over you as if she were making up the bed. What would it mean to say that she has *covered* you?

You can answer this question in two ways. One is to say that she has covered you provided every point on your body lies under some point of the quilt. That is, for every point of your body, there exists a point of the quilt that lies on top of the chosen point of your body. Another way to answer the question is similar, but expressed in terms of the individual pieces of cloth that were used in making the quilt.

Exercise 5.6.1 Use the universal and existential quantifiers to construct an alternate statement equivalent to your body being covered by the quilt.

Now suppose that the pieces of cloth used in the quilt are subsets of the real numbers, and think of the collection of these pieces as a family of sets. Stitching the pieces together to form the quilt is analogous to forming the union across the family of sets, forming a single set. Now imagine that your body is also a subset of the real numbers. Translating the quilt story into this set language, together with our first explanation of what it means to be covered by the quilt, motivates part of the following definition.

Definition 5.6.2 Suppose A is a set of real numbers and $\mathcal{C}$ is a family of sets of real numbers. If

$$A \subseteq \bigcup_{S \in \mathcal{C}} S \tag{5.12}$$

we say that $\mathcal{C}$ *covers* A, or is a *cover* of A. If every set in $\mathcal{C}$ is open, we call $\mathcal{C}$ an *open cover*. If $\mathcal{F}$ is a subfamily of $\mathcal{C}$ that also covers A, we call $\mathcal{F}$ a *subcover* of $\mathcal{C}$.

[20] Take a hint from the hint for Exercise 5.5.4.

Exercise 5.6.3 Translate your answer to Exercise 5.6.1 into appropriate mathematical language to state the definition of cover in Definition 5.6.2 in a slightly different way.

Example 5.6.4 Let $\mathcal{C} = \{N_1(x) : x \in \mathbb{Q}\}$, the set of neighborhoods of radius 1 of all rational numbers. Then $\mathcal{C}$ is an open cover of $\mathbb{R}$ because every real number is within 1 of some rational. Also, $\{N_1(n) : n \in \mathbb{Z}\}$ is a subcover of $\mathcal{C}$ because every real number is within 1 of some integer. ■

Exercise 5.6.5 What does it mean for $\mathcal{C}$ not to be a cover of A?

Covers and subcovers are of most interest to us when the cover $\mathcal{C}$ is an open cover with infinitely many sets in it, and $\mathcal{F}$ is a subcover that has only finitely many sets from $\mathcal{C}$ in it. In this case, we say that $\mathcal{C}$ can be reduced to a finite subcover. In the next exercises, part of what you will show is that an open cover of a set A cannot be reduced to a finite subcover. The natural way to do this is to show that any finite subfamily fails to cover A. These exercises will be very helpful when you prove part of the main result of this section.

Exercise 5.6.6 Let a be any real number. Then the family

$$\mathcal{C} = \{(-\infty, a - 1/n) \cup (a + 1/n, +\infty) : n \in \mathbb{N}\} \tag{5.13}$$

covers $\mathbb{R} - \{a\}$ but cannot be reduced to a finite subcover.

Exercise 5.6.7 The family of intervals $\mathcal{C} = \{(-n, n) : n \in \mathbb{N}\}$ covers $\mathbb{R}$ but has no finite subcover.

Now we are ready for the definition of compactness.

Definition 5.6.8 A set A is said to be *compact* if every open cover of A can be reduced to a finite subcover.

To show a given set is not compact, all we have to do is demonstrate an open cover that cannot be reduced to a finite subcover, as you did in Exercises 5.6.6 and 5.6.7. It might, therefore, be pretty easy to prove that entire classes of sets are not compact in this way. However, to demonstrate that a given set is compact from Definition 5.6.8 would appear to be quite a task, for something rather powerful must be shown about every conceivable open cover of the set. Thankfully, someone has been down this road before and supplied us with a relatively painless way to determine precisely which sets are compact.

Theorem 5.6.9 (Heine-Borel). A set is compact if and only if it is closed and bounded.

We will piece together the proof of Theorem 5.6.9 in stages by way of several exercises. The next two exercises comprise the $\Rightarrow$ direction of the proof.

Exercise 5.6.10 Suppose A is not closed. Demonstrate an open cover of A that has no finite subcover.[21]

Exercise 5.6.11 Suppose A is not bounded. Demonstrate an open cover of A that has no finite subcover.[22]

To prove the $\Leftarrow$ direction of Theorem 5.6.9, the following two theorems lay most of the groundwork. We will prove the first one here, for it is quite intricate. We leave the second as an exercise.

Theorem 5.6.12 The interval $[a, b]$ is compact.

Proof. Suppose $\mathcal{C}$ is an open cover of $[a, b]$. We construct a set E, which is by definition a subset of $(a, b]$, in the following way. For $x \in (a, b]$, let x be in E provided some finite subfamily of $\mathcal{C}$ covers the interval $[a, x]$. First we show that E is non-empty and bounded from above. Then we show that its LUB is b.

The fact that E is bounded from above is easy, for it is a subset of $(a, b]$ by definition, thus bounded from above by b. To show that E is non-empty, we note that since $\mathcal{C}$ covers $[a, b]$, then there is some S_1 that contains a. Since S_1 is open, there exists $\epsilon_1 > 0$ such that $N_{\epsilon_1}(a) \subseteq S_1$. Thus the subfamily $\{S_1\}$ covers $[a, a + \epsilon_1/2]$, so that $a + \epsilon_1/2 \in E$. Thus E is not empty.

Since E is a non-empty set bounded above by b, it has LUB $L \leq b$. First we show that $L \in E$. Then we show that $L = b$ in order to have that the entire interval $[a, b]$ can be covered by a finite subfamily of $\mathcal{C}$.

Since $\mathcal{C}$ is an open cover of $[a, b]$ and $L \leq b$, there exists $S_2 \in \mathcal{C}$ such that $L \in S_2$. Since S_2 is open, there exists $\epsilon_2 > 0$ such that $N_{\epsilon_2}(L) \subseteq S_2$. Since L is the LUB of E, there exists some $x \in (L - \epsilon_2, L] \cap E$. Thus $[a, x]$ can be covered by a finite subfamily $\mathcal{F}_2 \subseteq \mathcal{C}$. But then $\mathcal{F}_2 \cup \{S_2\}$ is a finite subfamily of $\mathcal{C}$ that covers $[a, L]$, so that $L \in E$.

Now suppose $L < b$. Since S_2 is open, there exists $\epsilon_3 > 0$ such that $N_{\epsilon_3}(L) \subseteq S_2$. Let $y = \min\{L + \epsilon_3/2, b\}$. Then $y > L$, $y \in S_2$, and $y \in [a, b]$. So the interval $[a, y]$ can be covered by a finite subcover of $\mathcal{C}$. This contradicts the fact that L is the LUB of E. Thus $L < b$ is impossible, and $L = b$. □

Knowing that $[a, b]$ is compact moves us very close to showing that any closed and bounded set is compact. The following exercise brings us within epsilon of being finished.

[21] Use the characterization of closed sets from Exercise 5.4.12 and Exercise 5.6.6.

[22] Use Exercise 5.6.7.

Exercise 5.6.13 If $A \subseteq B$, where A is closed and B is compact, then A is also compact.[23]

Now we have all the pieces we need to prove Theorem 5.6.9.

Proof of Theorem 5.6.9.

($\Rightarrow$) If A is not closed, then Exercise 5.6.10 implies that A is not compact. If A is not bounded, then Exercise 5.6.11 implies that A is not compact.

($\Leftarrow$) Suppose A is closed and bounded. Then there exists $M > 0$ such that $A \subseteq [-M, M]$. By Theorem 5.6.12, $[-M, M]$ is compact. Since A is a closed subset of a compact set, A is compact by Exercise 5.6.13. □

You might be a little disappointed that we went to all the trouble to define a new term (compactness) in complicated terms of open covers and subcovers, when it turns out that compact sets are precisely those that are closed and bounded, two ideas we already have a handle on. Well, you would be right to say there is no necessity to define a new term if the class of things to which the term applies can be easily described in already familiar language. But the point remains that closed and bounded sets have a very important feature: that every open cover can be reduced to a finite subcover. We need this feature in Section 7.7.

[23] If $\mathcal{C}$ is an open cover of A, then the fact that A^C is open should suggest a way to expand $\mathcal{C}$ into an open cover of B by the addition of one more set.

This page intentionally left blank

6

Sequences of Real Numbers

6.1 Sequences Defined

We can approach the subject of sequences in several ways. Let's move from an informal to a more formal way. A sequence of real numbers is an ordered list

$$\langle a_1, a_2, a_3, \ldots, a_n, \ldots\rangle \tag{6.1}$$

Listing the terms of a sequence in order might bring back memories of showing a set A is countable, for this listing is equivalent to constructing a function $f : \mathbb{N} \xrightarrow{\text{onto}} A$, where $f(1) = a_1$, $f(2) = a_2$, and so on. A *sequence* can therefore be defined as a function from the positive integers into the reals. The notation $\langle a_n\rangle_{n=1}^{\infty}$ is one standard way of denoting a sequence. Beginning the list with a_1 is a matter of convenience. There might be times when it seems more natural to begin with a_0.

If there is a formula for a_n, say, for example, $a_n = 1/n$, we have

$$\langle \tfrac{1}{n}\rangle_{n=1}^{\infty} = \langle 1, \tfrac{1}{2}, \tfrac{1}{3}, \ldots\rangle \tag{6.2}$$

We must distinguish (6.2) from $\{1/n\}_{n=1}^{\infty}$, which is merely the set of elements of the sequence $\langle 1/n\rangle$ and does not have the ordering of the elements as one of its defining characteristics. Saying $a_n = 1/n$ makes the idea of a sequence as a function more natural, for we are talking about nothing other than $f : \mathbb{N} \to \mathbb{R}$ defined by $f(n) = 1/n$, and the set $\{1/n\}$ can then be thought of as the range of the sequence. Then, if we want to visualize the sequence graphically, we can, as in Figure 6.1. Such a graph will help in visualizing limits in Section 6.2. Example 6.1.1 presents some examples of sequences we will refer to later.

Example 6.1.1

1. The sequence $\langle a_n\rangle_{n=1}^{\infty}$ defined by $a_n = \frac{1}{2} + (-1)^n \cdot \frac{1}{2}$ is the sequence $\langle 0, 1, 0, 1, 0, 1, \ldots\rangle$.
2. Letting $a_n = n + (-1)^n$ for $n \geq 0$ generates $\langle 1, 0, 3, 2, 5, 4, 7, 6, \ldots\rangle$.
3. If $a_n = \sin n\pi$, then we generate the sequence $\langle 0\rangle$.

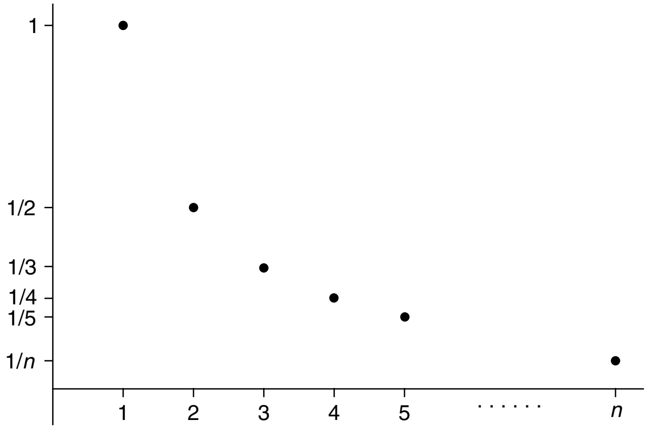

Figure 6.1 The sequence $a_n = f(n) = 1/n$.

4. We can define a_n in cases. For $n \geq 0$, let

$$a_n = \begin{cases} 0, & \text{if } n \text{ is even} \\ n^2 - 1 & \text{if } n \text{ is odd} \end{cases} \tag{6.3}$$

to generate the sequence $\langle 0, 0, 0, 8, 0, 24, 0, 48, 0, 80, 0, \ldots \rangle$.

5. If $a_n = n^2/2^n$ for $n \geq 1$, then we generate the sequence

$$\left\langle \tfrac{1}{2}, 1, \tfrac{9}{8}, 1, \tfrac{25}{32}, \tfrac{36}{64}, \tfrac{49}{128}, \tfrac{64}{256}, \ldots \right\rangle \qquad \blacksquare$$

6.1.1 Monotone Sequences

An important class of sequences derives from the following definition.

Definition 6.1.2 A sequence of real numbers $\langle a_n \rangle_{n=1}^{\infty}$ is said to be *increasing* (or *nondecreasing*) provided $a_m \leq a_n$ whenever $m < n$. It is said to be *decreasing* (or *nonincreasing*) provided $a_m \geq a_n$ whenever $m < n$. If $\langle a_n \rangle$ is either increasing or decreasing, then it is said to be *monotone*. If $a_m < a_n$ whenever $m < n$, then $\langle a_n \rangle$ is said to be *strictly increasing*. If $a_m > a_n$ whenever $m < n$, then $\langle a_n \rangle$ is said to be *strictly decreasing*.

If a and b are positive real numbers such that $a < b$, then Exercise 2.2.8 implies that $1/a > 1/b$. It follows that if $\langle a_n \rangle$ is an increasing (or decreasing) sequence of positive real numbers, then $\langle 1/a_n \rangle$ is a decreasing (or increasing) sequence of all positive numbers. The result is similar for sequences of negative real numbers. This suggests one way to approach the next example.

Example 6.1.3 Show that the sequence defined by $a_n = n/(n+1)$ for $n \geq 1$ is strictly increasing.

Solution We show that the sequence $\langle 1/a_n \rangle$ is strictly decreasing. Suppose $m < n$. Then $1/m > 1/n$, so that $1/a_m = 1 + 1/m > 1 + 1/n = 1/a_n$. Thus the sequence $\langle 1/a_n \rangle$ is strictly decreasing. ■

Given a positive real number x, the sequence $\langle x^n \rangle$ proves to be important in a number of different contexts.

Exercise 6.1.4 Let x be a positive real number. Then the sequence $\langle x^n \rangle_{n=0}^{\infty}$ is strictly increasing if $x > 1$ and strictly decreasing if $0 < x < 1$.[1]

Exercise 6.1.5 What does it mean to say that a sequence is not increasing?

Exercise 6.1.6 Give an example of a sequence that is both increasing and decreasing.

Exercise 6.1.7 Determine whether the following sequences are increasing, decreasing, or neither. Prove your claims.

(a) $\langle 1 - 1/n \rangle$

(b) $\langle 1/n^2 \rangle$

(c) $\langle 1 + (-1)^n/n^2 \rangle$

6.1.2 Bounded Sequences

Boundedness of sequences is defined in precisely the same way we defined boundedness of sets.

Definition 6.1.8 A sequence of real numbers $\langle a_n \rangle$ is said to be *bounded from above* provided there exists a real number M_1 such that $a_n \leq M_1$ for all n. Similarly, the sequence is said to be *bounded from below* provided there exists a real number M_2 such that $a_n \geq M_2$ for all n. The sequence is said to be *bounded* provided there exists $M > 0$ such that $|a_n| \leq M$ for all n. If a sequence is not bounded, it is said to be *unbounded*.

The boundedness terms in Definition 6.1.8 read very much like those in Definition 5.1.1. In fact, to say that a sequence is bounded from above, bounded from below, or bounded is precisely the same as saying that the range of the sequence is bounded from above, bounded from below, or bounded. Furthermore, we may apply Exercise 5.1.2 to the range of a sequence to have the following.

[1] See Exercise 3.5.7(b).

Theorem 6.1.9 A sequence is bounded if and only if it is bounded from above and below.

Example 6.1.10 From Example 6.1.1, we note without details that

1. $\langle \frac{1}{2} + (-1)^n \cdot \frac{1}{2} \rangle$ is bounded above by $M_1 = 8$ and below by $M_2 = 0$. It is therefore bounded.
2. $\langle n + (-1)^n \rangle$ is not bounded from above; therefore it is not bounded. It is, however, bounded from below by $M_2 = 0$.
3. The sequence $\langle a_n \rangle$ defined in Eq. (6.3) is not bounded from above, and thus is not bounded. It is bounded from below by $M_2 = -84$. ■

Example 6.1.11 We show that the sequence from Example 6.1.1 defined by $a_n = n^2/2^n$ is bounded. In Exercise 3.5.16, you showed that $2^n > n^2$ for all $n \geq 5$. Thus $0 < n^2/2^n < 1$ for all $n \geq 5$. Let $M = \max\{a_1, a_2, a_3, a_4, 1\}$. We show that $0 \leq a_n \leq M$ for all n.

Let n be a positive integer. If $1 \leq n \leq 4$, then clearly

$$0 \leq a_n \leq \max\{a_1, a_2, a_3, a_4\} \leq M \tag{6.4}$$

If $n \geq 5$, then $0 \leq a_n \leq 1 \leq M$. In either case, $0 \leq a_n \leq M$, so that the sequence is bounded. ■

EXERCISE 6.1.12 What does it mean for a sequence to be unbounded?

Example 6.1.13 Show that the sequence defined by $a_n = \sqrt{n}$ is unbounded.

Solution Pick $M > 0$, and let n be any integer satisfying $n > M^2$. By Exercise 3.9.4, we have that $a_n = \sqrt{n} > \sqrt{M^2} = M > 0$, so that $|a_n| > M$. ■

Since the sequence $\langle x^n \rangle$ is decreasing for $0 < x < 1$, it is bounded above by 1. Since every x^n is positive, $\langle x^n \rangle$ is bounded below by zero. Thus if $0 < x < 1$, the sequence $\langle x^n \rangle$ is bounded. However, the next example shows that $\langle x^n \rangle$ is unbounded if $x > 1$.

Example 6.1.14 Suppose $x > 1$. Then $\langle x^n \rangle_{n=0}^{\infty}$ is unbounded.

Solution Suppose $\langle x^n \rangle$ is bounded, and let L be the LUB of the sequence. Let $\epsilon = x - 1$, which is positive. Then by property M2, there exists a positive integer n such that $L - \epsilon < x^n \leq L$, so that $L < x^n + x - 1$. But since $x^n > 1$, we have that

$$x^{n+1} = x^{n+1} - x^n + x^n = x^n(x-1) + x^n > x - 1 + x^n > L \tag{6.5}$$

This is a contradiction, so that $\langle x^n \rangle$ is unbounded. ■

EXERCISE 6.1.15 Consider the sequence $\langle a_n \rangle$ defined by

$$a_n = \begin{cases} 10, & \text{if } n \text{ is odd} \\ 1 - n^2, & \text{if } n \text{ is even} \end{cases}$$

Prove that $\langle a_n \rangle$ is bounded from above but not from below.

To close out this introductory section on sequences, here is an example that illustrates a technique we will use in Section 6.2. The exercises that follow will be similar.

Example 6.1.16 Consider the sequence defined by $a_n = (n+3)/(n-1)$ for $n \geq 2$.

(a) Find a positive integer N such that $|a_N - 1| < 0.001$.

(b) Does the inequality in part (a) hold for all $n \geq N$?

Solution

(a) Expanding the inequality $|a_N - 1| < 0.001$ according to Theorem 2.3.4 and then solving for N yield the following equivalent inequality statements.

$$-0.001 < \frac{N+3}{N-1} - 1 < 0.001$$

$$-0.001 < \frac{4}{N-1} < 0.001$$

$$-0.001(N-1) < 4 < 0.001(N-1)$$

$$N > -3999 \quad \text{and} \quad N > 4001$$

The inequality $N > 4001$ is stronger than $N > -3999$, so we may let $N = 4002$.

(b) Suppose $n \geq N$. Then

$$-0.001 < 0 < \frac{4}{n-1} \leq \frac{4}{N-1} < 0.001 \tag{6.6}$$

so that $|a_n - 1| < 0.001$ for all $n \geq N$. ■

EXERCISE 6.1.17 Let $\langle a_n \rangle$ be defined as in Example 6.1.16, and let $\epsilon > 0$ be given. Find a value of N (which will be in terms of ϵ) such that $|a_N - 1| < \epsilon$. Show that $|a_n - 1| < \epsilon$ for all $n \geq N$.

Exercise 6.1.18 Let $\langle a_n \rangle$ be defined by $a_n = (n+3)/(2n-1)$, and suppose $\epsilon > 0$ is given. Find a value of N such that $|a_N - 1/2| < \epsilon$, and show that $|a_n - 1/2| < \epsilon$ for all $n \geq N$.

Exercise 6.1.19 Let $\langle a_n \rangle$ be defined by $a_n = n/(n+1)$, and suppose $0 < \epsilon < 1$ is given. Find a value of N such that $|a_N| > \epsilon$, and show that $|a_n| > \epsilon$ for all $n \geq N$.

Exercise 6.1.20 Let $\langle a_n \rangle$ be defined by $a_n = 1/(n^2+1)$, and suppose $\epsilon > 0$ is given. Find a value of N such that $|a_N| < \epsilon$, and show that $|a_n| < \epsilon$ for all $n \geq N$.[2]

Exercise 6.1.21 Suppose $\epsilon > 0$, and $\langle a_n \rangle$ and $\langle b_n \rangle$ are two sequences with the following properties. If $n \geq N_1$, then $|a_n - L_1| < \epsilon/2$, and if $n \geq N_2$, then $|b_n - L_2| < \epsilon/2$. Given this information, you should be able to make a true statement by filling in the blanks below.

$$\text{If } n \geq \underline{\quad ? \quad}, \text{ then } |(a_n + b_n) - (L_1 + L_2)| < \underline{\quad ? \quad}. \tag{6.7}$$

Fill in the blanks with the appropriate numbers, then prove that the resulting statement is true.

6.2 Convergence of Sequences

One of the most important questions we can ask about a sequence of real numbers is whether it *converges*. If a sequence converges to a real number, its terms are "homing in on" or "approaching" some real number as n takes on larger and larger values. We can also talk about a sequence converging to positive or negative infinity. In this section, we define these notions and prove some important basic results about convergent sequences.

It would be misleading not to say up front that the idea of convergence created no small amount of controversy and distress in the historical development of mathematics. It took quite some time for a satisfactory definition to be devised. Sir Isaac Newton, in developing the calculus, assumed some pretty sweeping results concerning convergence in his work, only to have its details placed on a firm footing later. Because the idea is a little complicated, we will work our way into it a little at a time.

6.2.1 Convergence to a Real Number

Figure 6.2 illustrates convergence to a real number L. Given some $\epsilon > 0$, the ϵ-neighborhood of L contains all the terms in the sequence beyond some threshold term a_N. In order to visualize how you might prove that the terms of a sequence $\langle a_n \rangle$ are converging to L, imagine the following task. Someone gives you an $\epsilon > 0$.

[2] You will have to treat the cases $\epsilon > 1$ and $\epsilon \leq 1$ separately.

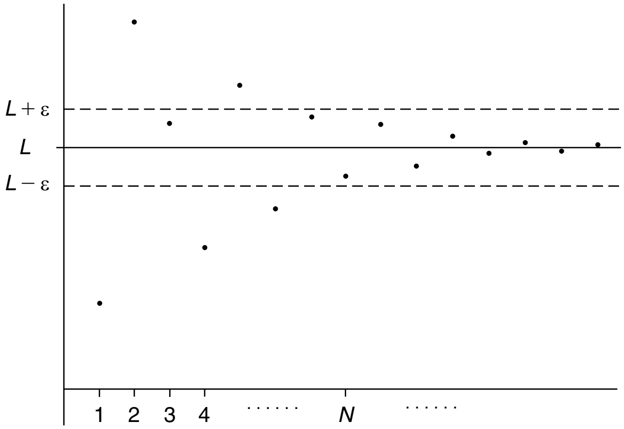

Figure 6.2 A convergent sequence.

You then take that ϵ value and do a quick calculation to determine some positive integer N (which will depend on ϵ) with the property that the Nth term and all those after it fall in the ϵ-neighborhood of L. If someone else then gives you an even smaller $\epsilon > 0$, you can still find a threshold value N with the same property, though it will probably be a higher threshold. If you are given any $\epsilon > 0$ and are able to find some positive integer N with the property that all terms from the Nth one onward fall in the ϵ-neighborhood of L, then you will have shown that the sequence converges to L.

Before we give the definition of convergence, note that you have already done some of this sort of work in the exercises of Section 6.1. In Exercise 6.1.17, you showed that

$$\text{If } n > 1 + \frac{4}{\epsilon}, \text{ then } \left| \frac{n+3}{n-1} - 1 \right| < \epsilon.$$

That is, all terms with an index larger than $1 + 4/\epsilon$ fall in the ϵ-neighborhood of 1. In Exercise 6.1.18, you showed that

$$\text{If } n > \frac{1}{2} + \frac{7}{4\epsilon}, \text{ then } \left| \frac{n+3}{2n-1} - \frac{1}{2} \right| < \epsilon.$$

That is, all terms with an index larger than $1/2 + 7/4\epsilon$ fall in the ϵ-neighborhood of $1/2$. Note that in both of these exercises, smaller values of ϵ yield higher thresholds for n. Now we are ready for the definition of convergence.

Definition 6.2.1 Suppose $\langle a_n \rangle_{n=1}^{\infty}$ is a sequence of real numbers. We say that the sequence *converges* if there exists a real number L such that, for every $\epsilon > 0$, there exists a positive integer N with the property that $n \geq N$ implies $|a_n - L| < \epsilon$. That is, the sequence converges provided

$$(\exists L \in \mathbb{R})(\forall \epsilon > 0)(\exists N \in \mathbb{N})(\forall n \in \mathbb{N})(n \geq N \to |a_n - L| < \epsilon) \tag{6.8}$$

If $\langle a_n \rangle_{n=1}^{\infty}$ converges to L, we say that L is the *limit* of the sequence as n approaches infinity, and we write this as

$$\lim_{n\to\infty} a_n = L \tag{6.9}$$

We also say that a_n approaches L as n approaches infinity, and we write this as

$$a_n \to L \text{ as } n \to \infty \tag{6.10}$$

If a sequence does not converge, we say it *diverges*.

To construct a proof of convergence is, as always, to follow the logical flow of the definition. A general form for such a proof would look something like this.

Theorem 6.2.2 (Sample). The sequence $\langle a_n \rangle_{n=1}^{\infty}$ converges to L.

Proof. Let $\epsilon > 0$ be given. Let N be any integer satisfying $N > f(\epsilon)$.[3] Then if $n \geq N$, we have that

$$|a_n - L| = f^{-1}(n) \leq f^{-1}(N) < \epsilon \tag{6.11}$$

Thus $a_n \to L$ as $n \to \infty$. □

Here is a cleaned up example.

Example 6.2.3 The sequence $\langle (n+1)/(4n+3) \rangle$ converges to $1/4$.

Solution Let $\epsilon > 0$ be given. Let N be any positive integer satisfying $N > 1/16\epsilon - 3/4$. Then if $n \geq N$, we have

$$\left| \frac{n+1}{4n+3} - \frac{1}{4} \right| = \left| \frac{1}{16n+12} \right| < \left| \frac{1}{16N+12} \right| < \epsilon \tag{6.12}$$

Thus $(n+1)/(4n+3) \to 1/4$ as $n \to \infty$. ■

If you wonder what inspired the statement $N > 1/16\epsilon - 3/4$, it is the result of working backward from the last step of the inequality in (6.12) and solving for N to produce a logically equivalent inequality. This scratchwork is what makes inequality (6.12) work.

Notice that the logical flow of Theorem 6.2.2 and Example 6.2.3 does not exploit the expression $(\exists L \in \mathbb{R})$ given in Definition 6.2.1 because the supposed limit is provided in the statement of the theorem. This is typical of convergence proofs of this sort, because your theorem will almost certainly come with the limit as part of the package.

[3] Where you have already done the scratchwork to find what $f(\epsilon)$ should be.

Exercise 6.2.4 Write polished proofs that the following sequences from Section 6.1 converge to the given limit.

(a) $(n+3)/(n-1) \to 1$ (Exercise 6.1.17)

(b) $1/(n^2+1) \to 0$ (Exercise 6.1.20)

Exercise 6.2.5 By stripping off the $(\exists L \in \mathbb{R})$ in the definition of convergence, state what it means for a sequence $\langle a_n \rangle$ not to converge to a given real number L as $n \to \infty$.

Specific examples of convergence are great, but we need generalized theorems addressing convergence in order to develop a broader theory of convergence. The following results go a long way by giving us some basic building blocks and ways of combining them.

Exercise 6.2.6 The limit of a sequence of real numbers is unique.[4]

Theorem 6.2.7 Let c be a real number. Then the constant sequence defined by $a_n = c$ converges to c. That is, $\lim_{n\to\infty} c = c$.

Proof. Let $\epsilon > 0$ be given. Let $N = 1$. Then if $n \geq N$, we have that

$$|a_n - L| = |c - c| = 0 < \epsilon$$

□

Exercise 6.2.8 The sequence $1/n \to 0$ as $n \to \infty$.[5]

Exercise 6.2.9 Show that $1/2^n \to 0$ as $n \to \infty$.[6]

The proof of the next theorem reveals a useful trick in showing how a sequence converges, so take note. We exploit the convergence of $\langle a_n \rangle$ by applying its defining feature to ϵ/M.

Theorem 6.2.10 Suppose $a_n \to 0$ and $\langle b_n \rangle$ is bounded. Then $a_n b_n \to 0$ as $n \to \infty$.

Proof. Let $\epsilon > 0$ be given. Since $\langle b_n \rangle$ is bounded, there exists $M > 0$ such that $|b_n| \leq M$ for all n. Now ϵ/M is also positive. Since $a_n \to 0$, there exists a positive integer N such that $|a_n| < \epsilon/M$ for all $n \geq N$. Thus if $n \geq N$, we have

$$|a_n b_n - 0| = |a_n b_n| = |a_n|\,|b_n| < \frac{\epsilon}{M} \cdot M = \epsilon \tag{6.13}$$

Thus $a_n b_n \to 0$ as $n \to \infty$. □

[4] If L_1 and L_2 are both limits and $L_1 < L_2$, then letting $\epsilon = (L_2 - L_1)/2$ should yield a contradiction.

[5] Use the Archimedean property.

[6] See Exercise 3.5.16.

Convergence is stronger than boundedness, but in one part of upcoming Exercise 6.2.12, all you need to know about a convergent sequence is that it is bounded. To prove the result in the next exercise, consider the following suggestion. If $\langle a_n \rangle$ converges to L, its terms eventually settle down close to L, so that all the terms beyond a certain point are caught in the interval $(L-1, L+1)$. So no matter how wildly the sequence might jump around before this point, there are only finitely many terms to contain. This should suggest a value of M to serve as a bound.

EXERCISE 6.2.11 A convergent sequence is bounded.[7,8]

EXERCISE 6.2.12 Suppose $a_n \to L_1$ and $b_n \to L_2$ as $n \to \infty$. Then

(a) $\lim_{n\to\infty}(a_n + b_n) = L_1 + L_2$[9]

(b) $\lim_{n\to\infty}(a_n b_n) = L_1 L_2$[10,11]

Corollary 6.2.13 Suppose $\lim_{n\to\infty} a_n = L$ and let c be any real number. Then $\lim_{n\to\infty}(ca_n) = cL$.

Proof. Let $b_n = c$ for all n in Exercise 6.2.12(b) and apply Theorem 6.2.7. □

Corollary 6.2.14 If $a_n \to L_1$ and $b_n \to L_2$ as $n \to \infty$, then $(a_n - b_n) \to (L_1 - L_2)$ as $n \to \infty$.

Proof.

$$\begin{aligned}\lim_{n\to\infty}(a_n - b_n) &= \lim_{n\to\infty}[a_n + (-1)b_n] = \lim_{n\to\infty} a_n + \lim_{n\to\infty}(-1)b_n \\ &= L_1 + (-1)L_2 = L_1 - L_2\end{aligned} \tag{6.14}$$

□

With an induction argument calling on Exercises 6.2.8 and 6.2.12(b), you can show the following.

EXERCISE 6.2.15 If p is any positive integer, then $1/n^p \to 0$ as $n \to \infty$.

[7] Look at the reasoning we used in Example 6.1.11 where we took care of finitely many terms individually and then handled all the rest with a single number. A good start for this proof would be to say: "Let $\epsilon = 1$. Then there exists . . .."

[8] Let $M = \max\{|a_1|, |a_2|, \ldots, |a_{N-1}|, |L|+1\}$.

[9] Remember, if $\epsilon > 0$, then so is $\epsilon/2$. See Exercise 6.1.21.

[10] First deal with the case $L_1 = 0$ by appealing to previous results. For $L_1 \neq 0$, do some preliminary playing around with the expression $|a_n b_n - L_1 L_2|$. The sneaky trick here is adding and subtracting the same quantity inside $|a_n b_n - L_1 L_2|$. If you want to know what that quantity is, check the next hint.

[11] Add and subtract $b_n L_1$ and use the triangle inequality. Also, apply Exercise 6.2.11 to $\langle b_n \rangle$.

To combine two sequences $\langle a_n \rangle$ and $\langle b_n \rangle$ by division to produce $\langle a_n/b_n \rangle$, we first need to do a little work on the sequence $\langle 1/b_n \rangle$ by itself. We would like to be able to say something to the effect that the limit of a quotient is the quotient of the limits. Well, we can, sort of.

The sequence $\langle 1/n \rangle$ contains only positive terms. However, $1/n \to 0$, so every ϵ-neighborhood of zero contains some $1/n$. Thus the terms of the sequence $\langle 1/n \rangle$ cannot be bounded away from zero. But if the terms of a sequence are all nonzero and converge to a nonzero limit, then we can bound its terms away from zero.

Exercise 6.2.16 Suppose $\langle b_n \rangle$ is a sequence of *nonzero* real numbers such that $b_n \to L$, where L is nonzero. Then there exists $M > 0$ such that $|b_n| \geq M$ for all n.[12,13]

Exercise 6.2.16 provides the right machinery for you to prove the following.

Exercise 6.2.17 Suppose $\langle b_n \rangle$ is a sequence of nonzero real numbers that converges to a nonzero limit L. Then $\lim_{n\to\infty} 1/b_n = 1/L$.

Corollary 6.2.18 Suppose $\langle a_n \rangle$ is a sequence that converges to L_1, and $\langle b_n \rangle$ is a sequence of nonzero real numbers that converges to a nonzero limit L_2. Then $a_n/b_n \to L_1/L_2$ as $n \to \infty$.

Proof. Apply Exercise 6.2.12(b) to $\langle a_n \rangle$ and $\langle 1/b_n \rangle$. □

Many of our results from Theorem 6.2.7 through Corollary 6.2.18 can take us a long way in dealing with a particular class of sequences. We will prove one case of the next theorem and leave the other case to you as an exercise.

Theorem 6.2.19 Suppose $\langle a_n \rangle$ is a sequence defined by

$$a_n = \frac{b_r n^r + b_{r-1} n^{r-1} + \cdots + b_1 n + b_0}{c_s n^s + c_{s-1} n^{s-1} + \cdots + c_1 n + c_0} \tag{6.15}$$

where all b_k and c_k are real numbers, b_r and c_s are nonzero, and the denominator is nonzero for all n. If $r < s$, then $a_n \to 0$ as $n \to \infty$. If $r = s$, then $a_n \to b_r/c_s$ as $n \to \infty$.

Proof of case r<s. Suppose $r < s$. Multiply the expression for a_n in Eq. (6.15) by n^{-s}/n^{-s} to have

$$a_n = \frac{b_r n^{r-s} + b_{r-1} n^{r-1-s} + \cdots + b_1 n^{1-s} + b_0 n^{-s}}{c_s + c_{s-1} n^{-1} + \cdots + c_1 n^{1-s} + c_0 n^{-s}} \tag{6.16}$$

[12] First deal with $L > 0$. Let $\epsilon = L/2$. Then there exists a positive integer N such that . . ..

[13] If $L < 0$, then use the fact that $-b_n \to -L$.

Since $r < s$, every exponent of n in Eq. (6.16) is negative, and we may apply the results of this section to have

$$\begin{aligned}
\lim_{n\to\infty} a_n &= \frac{\lim(b_r n^{r-s} + \cdots + b_0 n^{-s})}{\lim(c_s + \cdots + c_0 n^{-s})} \\
&= \frac{\lim(b_r n^{r-s}) + \cdots + \lim(b_0 n^{-s})}{\lim c_s + \cdots + \lim(c_0 n^{-s})} \\
&= \frac{(\lim b_r)(\lim n^{r-s}) + \cdots + (\lim b_0)(\lim n^{-s})}{\lim c_s + \cdots + (\lim c_0)(\lim n^{-s})} \\
&= \frac{b_r \cdot 0 + \cdots + b_0 \cdot 0}{c_s + \cdots + c_0 \cdot 0} = \frac{0}{c_s} = 0
\end{aligned} \tag{6.17}$$

□

EXERCISE 6.2.20 Prove the case of Theorem 6.2.19 where $r = s$.

EXERCISE 6.2.21 Suppose $a_n \to L_1$ and $b_n \to L_2$ as $n \to \infty$. Suppose also that $a_n \leq b_n$ for all n. Show that $L_1 \leq L_2$.[14]

EXERCISE 6.2.22 Show that if $a_n \to L$ as $n \to \infty$, then $|a_n| \to |L|$ as $n \to \infty$.[15]

EXERCISE 6.2.23 The Sandwich Theorem: If $a_n \leq c_n \leq b_n$ for all n, and if $a_n \to L$ and $b_n \to L$, then $c_n \to L$.

6.2.2 Convergence to Infinity

Sometimes a sequence does not converge to a real number, but it does exhibit a nice behavior in that it increases (or decreases) without bound. The term we use for such behavior is that $a_n \to +\infty$ (or $-\infty$). We have already used the somewhat loose expression $n \to \infty$ in Definition 6.2.1, even though infinity is not a real number. It is not difficult to define the notion of $\lim_{n\to\infty} a_n = \infty$, but there is no way to do it in terms of neighborhoods, unless you can concoct some sort of definition of a neighborhood of infinity. We will do exactly that in Section 7.4, where we will spend some quality time with infinity. For now, we will be content to work with a few basic ideas.

If $a_n \to L$ as $n \to \infty$, the terms might hop all around L, just as long as that hopping stays within an ϵ-radius of L beyond some threshold term. That this bouncing around can be eventually contained for all $\epsilon > 0$ is why the bouncing apparently fades out. To say $a_n \to +\infty$ as $n \to \infty$ means loosely that the terms a_n are hopping

[14] Try proof by contrapositive.

[15] See Exercise 2.3.10.

increasingly higher and higher, even though there is no reason the terms cannot do at least a little hopping back down. Instead of selecting an arbitrary ϵ-radius around some L, we set an arbitrary hurdle bar at some positive number M and insist that all terms beyond some threshold are above the hurdle. If every M is associated with some threshold beyond which $a_n \geq M$, then we say $a_n \to +\infty$ as $n \to \infty$. Here is the definition.

Definition 6.2.24 Suppose $\langle a_n \rangle$ is a sequence of real numbers. We say that $\lim_{n\to\infty} a_n = +\infty$, provided for all $M > 0$, there exists a positive integer N with the property that $n \geq N$ implies $a_n > M$. Logically, we may write this as

$$(\forall M > 0)(\exists N \in \mathbb{N})(\forall n \in \mathbb{N})(n \geq N \to a_n > M) \tag{6.18}$$

We also write $a_n \to +\infty$ as $n \to \infty$, and we sometimes say that the sequence *increases without bound.*

Example 6.2.25 Show that $\lim_{n\to\infty} \sqrt{n} = +\infty$.

Solution Pick $M > 0$ and let N be any positive integer satisfying $N > M^2$. Then if $n \geq N$, we have that $\sqrt{n} \geq \sqrt{N} > \sqrt{M^2} = M$. ■

Exercise 6.2.26 If $\langle a_n \rangle$ is increasing and unbounded, then $a_n \to +\infty$ as $n \to \infty$.

An immediate consequence of Exercise 6.2.26 is that if $x > 1$, then $x^n \to +\infty$ as $n \to \infty$. For Exercise 6.1.4 showed that $\langle x^n \rangle$ is increasing for $x > 1$, and in Example 6.1.14 we showed it is unbounded.

Exercise 6.2.27 Suppose $\langle a_n \rangle$ is a sequence of nonzero real numbers such that $a_n \to +\infty$ as $n \to \infty$. Then $1/a_n \to 0$ as $n \to \infty$.

Exercise 6.2.28 If $0 < x < 1$, then $x^n \to 0$ as $n \to \infty$.

Exercise 6.2.29 If $-1 < x < 0$, then $x^n \to 0$ as $n \to \infty$.

Exercise 6.2.30 In the spirit of Definition 6.2.24, create a definition for the statement $\lim_{n\to\infty} a_n = -\infty$.

6.3 The Nested Interval Property

Imagine a sequence of closed intervals $\langle [a_n, b_n] \rangle_{n=1}^{\infty}$, where $[a_n, b_n] \subseteq [a_m, b_m]$ whenever $m < n$. Such a sequence of intervals is said to be *nested* because any particular interval contains all those after it. In this section, we derive the

Nested Interval Property (NIP) of the real numbers, which says that a nested sequence of closed intervals whose lengths $(b_n - a_n)$ tend to zero as $n \to \infty$ will intersect down to a single real number. We will then investigate an important implication of the NIP as it relates to bounded sequences. We also demonstrate that the NIP can be used axiomatically to prove the LUB property, thus making the two statements logically equivalent.

6.3.1 From LUB Axiom to NIP

If $\langle [a_n, b_n] \rangle$ is a nested sequence of closed intervals, then clearly $a_m \le a_n < b_n \le b_m$ whenever $m < n$. In particular, $a_1 \le a_n < b_n \le b_1$ for all n. Thus the sequence of left end points $\langle a_n \rangle$ is monotone increasing and bounded above by b_1, and the sequence of right end points $\langle b_n \rangle$ is monotone decreasing and bounded below by a_1. As we begin to make our way to the NIP, we need the following.

Exercise 6.3.1 Suppose $\langle a_n \rangle$ is an increasing sequence of real numbers that is bounded from above. Then $\langle a_n \rangle$ converges to its LUB.

Exercise 6.3.1 is the step in the process of proving the NIP where you will use the LUB axiom. The next result should follow quickly from Exercise 6.3.1 by the strategic insertion of some negative signs.

Exercise 6.3.2 Suppose $\langle b_n \rangle$ is a decreasing sequence of real numbers that is bounded from below. Then $\langle b_n \rangle$ converge to its GLB.

If the hypothesis conditions for Exercise 6.3.1 are satisfied, we say that $\langle a_n \rangle$ converges to *L from below*, and we write $a_n \nearrow L$. Similarly for Exercise 6.3.2, we say $\langle b_n \rangle$ converges to *L from above*, and we write $b_n \searrow L$.

Thus we have that a nested sequence of closed intervals $\langle [a_n, b_n] \rangle$ has the property that $a_n \to L_1$ and $b_n \to L_2$ for some real numbers L_1 and L_2. Now if $\langle [a_n, b_n] \rangle$ also has the property that the lengths of the intervals approach zero as $n \to \infty$, then we have $\lim_{n\to\infty}(b_n - a_n) = 0$. Because $\langle a_n \rangle$ and $\langle b_n \rangle$ converge individually, then by Corollary 6.2.14,

$$0 = \lim_{n\to\infty}(b_n - a_n) = \lim_{n\to\infty} b_n - \lim_{n\to\infty} a_n = L_2 - L_1 \tag{6.19}$$

Thus $L_1 = L_2$. The last piece of the NIP puzzle asserts that this limit is the single element in the intersection of all the nested intervals in the sequence. Here is a statement of the theorem and its proof, with its last detail left to you as an exercise.

Theorem 6.3.3 (Nested Interval Property). Let $\langle [a_n, b_n] \rangle$ be a sequence of nested intervals with the property that $\lim_{n\to\infty}(b_n - a_n) = 0$. Then $\cap_{n=1}^{\infty}[a_n, b_n]$ contains precisely one real number.

Proof. By Exercises 6.3.1 and 6.3.2, both $\langle a_n \rangle$ and $\langle b_n \rangle$ converge, and by Corollary 6.2.14, they have the same limit L. We show that $\cap_{n=1}^{\infty}[a_n, b_n] = \{L\} \ldots$ □

EXERCISE 6.3.4 Finish the proof of Theorem 6.3.3.

6.3.2 The NIP Applied to Subsequences

In Section 6.4, we will need some results based on taking a given sequence and creating from it a new sequence by perhaps deleting some of its terms and preserving the order of the kept terms. The new sequence is called a *subsequence*, and we take the opportunity now to introduce the term and note what the NIP has to say about certain sequences and subsequences.

Creating a subsequence of a given sequence can make for a little notational mess when it comes to indexing the subsequence, so let's look at an example. Given a sequence $\langle a_n \rangle$, suppose we want to consider the following subsequence.

$$\langle a_2, a_8, a_{10}, a_{40}, a_{44}, a_{66}, \ldots \rangle \tag{6.20}$$

The notational dilemma can be resolved by noting that the indices

$$\langle 2, 8, 10, 40, 44, 66, \ldots \rangle \tag{6.21}$$

form a strictly increasing sequence $\langle n_k \rangle_{k=1}^{\infty}$ of positive integers. If we denote $n_1 = 2$, $n_2 = 8$, $n_3 = 10$, and so on, then the sequence $\langle n_1, n_2, n_3, \ldots \rangle$ is a way of indexing our subsequence. The subsequence can then be denoted by

$$\langle a_{n_1}, a_{n_2}, a_{n_3}, a_{n_4}, \ldots \rangle = \langle a_{n_k} \rangle_{k=1}^{\infty} \tag{6.22}$$

where k is now the indexing variable. This brings us to a definition.

Definition 6.3.5 Suppose $\langle a_n \rangle_{n=1}^{\infty}$ is a sequence of real numbers, and suppose $\langle n_k \rangle_{k=1}^{\infty}$ is a strictly increasing sequence of positive integers. Then the sequence $\langle a_{n_k} \rangle_{k=1}^{\infty}$ is called a *subsequence* of the sequence $\langle a_n \rangle$.

Suppose you love to play darts. You play every day from now to eternity, and you're pretty good at it. At least you're good enough so that you always hit somewhere on the target. If every dart you throw leaves a tiny hole behind, what must eventually happen to the set of pinholes on the target as time goes on? Well, it is certainly possible that your darts could always land on one of a finite number of spots on the target, in which case you would hit at least one spot infinitely many times. But if you do not hit any one spot infinitely many times, then your target will have infinitely many holes in it. If so, then the finite size of the target will imply that the pinholes are going to cluster in one or more places. Of course, if we think of two direct hits on the same exact spot as leaving two distinguishable holes (why not?), then it can indisputably be said of your target that the holes are *somewhere* clustered and around at least one spot.

This little analogy is going somewhere. Suppose $\langle a_n \rangle$ is a sequence of real numbers that is bounded by $M > 0$. Then the interval $[-M, M]$ is the dart board target, and the terms of the sequence are the spots you hit. The next theorem says that there are an infinitely many terms of the sequence clustered around some point in $[-M, M]$, but it does so in the language of subsequences. It will come in really handy in Exercise 6.3.7, where you prove the famous Bolzano-Weierstrass Theorem.

Theorem 6.3.6 Every bounded sequence of real numbers contains a convergent subsequence.

Proof. Suppose $\langle a_n \rangle$ is bounded by M, and denote $[c_0, d_0] = [-M, M]$. (See Fig. 6.3.) Let m_0 be the midpoint of $[c_0, d_0]$, and note that either $[c_0, m_0]$ or $[m_0, d_0]$ must contain infinitely many of the a_n. (Perhaps both do.) Let $[c_1, d_1]$ be one of the subintervals of $[c_0, d_0]$ that contains infinitely many of the a_n, pick any term of the sequence from $[c_1, d_1]$, and denote it as a_{n_1}. Repeat this process by letting m_1 be the midpoint of $[c_1, d_1]$, letting $[c_2, d_2]$ be whichever of $[c_1, m_1]$ or $[m_1, d_1]$ contains infinitely many of the a_n, and then picking any term of the sequence from $[c_2, d_2]$ whose index exceeds n_1, and call this term a_{n_2}. Continuing in this fashion, we generate two things: a nested sequence of intervals $[c_k, d_k]$ whose lengths are

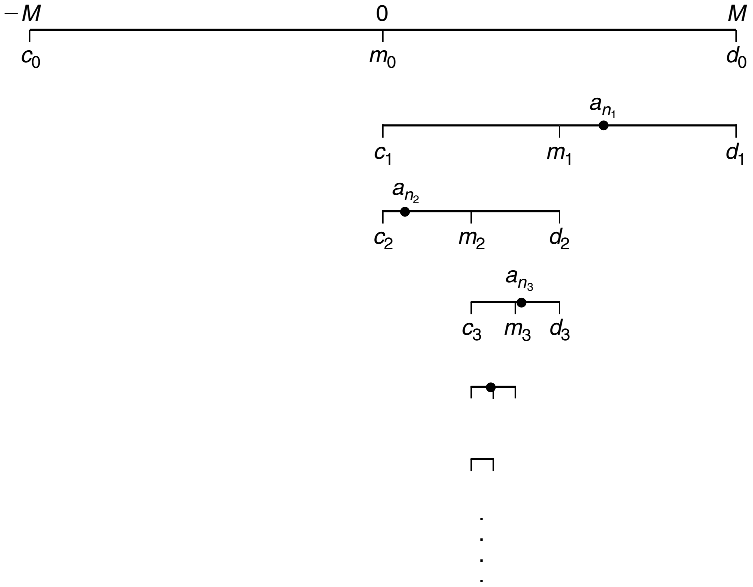

Figure 6.3 Nested intervals generated in the proof of Theorem 6.3.6.

$2M/2^k$ and therefore tend to zero (Exercise 6.2.9), and a subsequence $\langle a_{n_k} \rangle_{k=1}^{\infty}$, each term of which falls in $[c_k, d_k]$.

By the NIP, $\cap_{k=1}^{\infty}[c_k, d_k] = \{x_0\}$ for some real number x_0. Since $c_k \to x_0$ and $d_k \to x_0$ as $k \to \infty$, and since $c_k \leq a_{n_k} \leq d_k$ for all k, the Sandwich Theorem (Exercise 6.2.23) implies that $a_{n_k} \to x_0$. □

EXERCISE 6.3.7 Prove the Bolzano-Weierstrass Theorem: Every bounded, infinite set of real numbers has a cluster point.[16]

6.3.3 From NIP to LUB Axiom

Let's go back and note how the LUB axiom has contributed to the results we have developed thus far. We said in Section 3.10 that assumption A22 can be derived from the LUB axiom, though we have not proved this is true. All the results of Sections 3.9 and 3.10 are based on A22, and hence depend on the LUB axiom as well. Certainly, a lot of our work in Chapter 5 relied on the LUB axiom. Section 5.1 was all about the LUB axiom. Theorem 5.3.14 used the LUB axiom to show that the empty set and the real numbers are the only subsets of the real numbers that are both open and closed. In Section 6.1, the only places we used any results from the LUB axiom were in choosing the maximum value of a finite set in proving some results about boundedness. Consequently, the theorems from Section 6.2 involving boundedness stem from the LUB axiom.

We now want to delete the LUB axiom from our list of assumptions and replace it with the NIP. We can then show as a theorem that the set of real numbers has the LUB property. To do this, we retain our terminology of sequences and convergence. Assuming the NIP, the next exercise is the result we want to prove. If you want more elaboration on the task and some suggestions for how to tackle it, read the comments that follow.

EXERCISE 6.3.8 Every non-empty set of real numbers that is bounded from above has a least upper bound in the real numbers.[17]

Your task is to use the NIP to prove the LUB property as a theorem. That is, somehow you are going to use what you know about the set to generate a sequence of nested intervals whose lengths tend to zero, so that you can apply the NIP to this sequence. If you want more guidance, read on for a thumbnail sketch of the proof.

Suppose A is a non-empty set bounded above by M. Since A is non-empty, we can choose some $x \in A$. We can then split the interval $[x, M]$ in half repeatedly. At

[16] You can apply Theorem 6.3.6 by creating just the right sequence from the set, or you can mimic the proof of Theorem 6.3.6 in its use of the NIP.

[17] Prove that the single element of $\cap[a_n, b_n]$ has properties L1–L2 by contradiction.

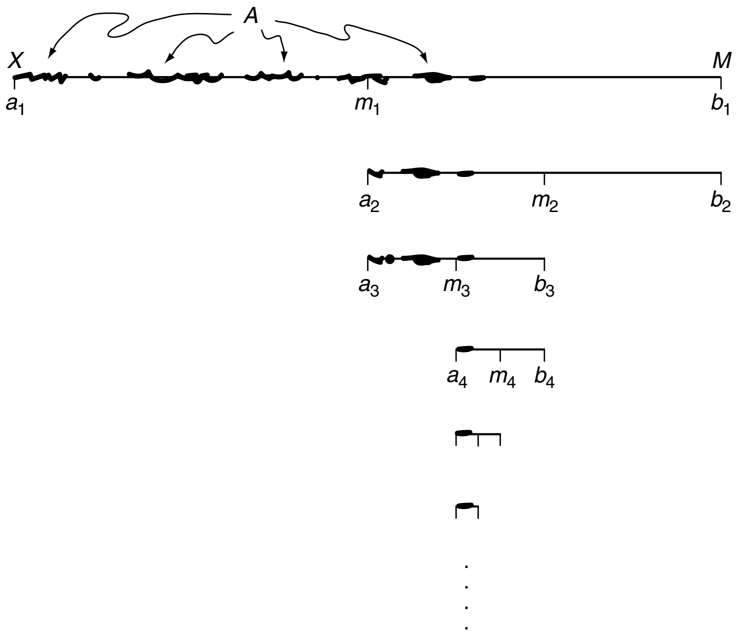

Figure 6.4 Nested intervals generated in the proof of Exercise 6.3.8.

each stage, keep the right half of the split interval if it contains any elements of A; otherwise, keep the left half. (See Figure 6.4.) The successive splits will produce a sequence of nested intervals whose lengths tend to zero. The NIP then gives us a single real number, the intersection of all the intervals, which you can show has properties L1–L2 (or M1 and M2 if you prefer).

6.4 Cauchy Sequences

Now we turn our attention to a certain class of sequences called *Cauchy* sequences. If we can say that convergent sequences are characterized by the fact that the terms all eventually bunch up around some real number (the limit), then we say that Cauchy sequences are characterized by the fact that the terms all eventually bunch up around each other, without any specific reference to a possible real number limit. This bunching of terms is logically equivalent to convergence to a real number limit (with the help of the NIP), and this is the main result we want to prove in this section. Once we work through these results, we will discuss why Cauchy sequences are important in more general mathematical settings, and we relate their convergence to the LUB axiom and the NIP.

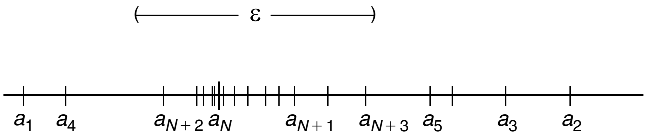

Figure 6.5 A Cauchy sequence.

6.4.1 Convergence of Cauchy Sequences

First let's define what we mean for a sequence to be Cauchy. Then we will show that a sequence is convergent if and only if it is Cauchy.

Definition 6.4.1 A sequence of real numbers $\langle a_n \rangle$ is said to be *Cauchy* provided that, for all $\epsilon > 0$, there exists a positive integer N such that $|a_m - a_n| < \epsilon$ for all $m, n \geq N$.

One way to envision a Cauchy sequence in terms of some given $\epsilon > 0$ and the corresponding threshold N is that all terms of index at least N fall within an ϵ distance, not only of a_N, but of each other as well. See Figure 6.5. There is no apparent limit around which the neighborhood is anchored, so the question of whether a Cauchy sequence must converge is not immediately answerable. In showing that convergence and Cauchy are equivalent, let's do the easy part first. The next exercise says merely that if the terms of a sequence bunch up around a real number limit, then they must bunch up around each other.

Exercise 6.4.2 A convergent sequence is Cauchy.

Exercise 6.4.3 What does it mean for a sequence $\langle a_n \rangle$ not to be Cauchy?

If a sequence is not Cauchy, then by Exercise 6.4.2 it does not converge. Sometimes a natural way to show that a particular sequence diverges is to show it is not Cauchy.

Exercise 6.4.4 Show that the sequence $\langle (-1)^n \rangle$ is not Cauchy and hence diverges.

To prove that a Cauchy sequence is convergent, we need the following lemma. To prove it, mimic your proof of Exercise 6.2.11. Only now, instead of fencing in the tail of the sequence around a known limit, bound it around the threshold term a_N by using $\epsilon = 1$.

EXERCISE 6.4.5 A Cauchy sequence is bounded.

At this point, you have all the machinery you need to show that a Cauchy sequence is convergent. Here is a statement of the theorem, with some suggestions for the proof following.

EXERCISE 6.4.6 A Cauchy sequence of real numbers is convergent.

If a sequence is Cauchy, then by Exercise 6.4.5 it is bounded. Thus it contains a convergent subsequence by Theorem 6.3.6. This gives you a real number limit to work with, and you can show that the entire sequence converges to this same limit by using convergence of the subsequence and Cauchiness of the sequence.

The real numbers are only one example of a set where we address the convergence of Cauchy sequences. In any set endowed with a *metric* (a measure of distance between elements), Cauchy sequences can be defined in the same way as in Definition 6.4.1 by interpreting $|a_m - a_n|$ as the distance between elements in that context. Any such *metric space*, as we call it, where Cauchy sequences always converge to an element in the space is said to be *complete*.

In some situations, completeness might be taken as an axiom of the metric space. What we have done here is to show completeness of the real numbers as a theorem, based on the NIP, hence based on the LUB axiom. Some authors of texts in analysis prefer to go the other way by assuming completeness as an axiom of the real numbers and demonstrating the NIP and LUB property as theorems. Of course, this means that completeness is logically equivalent to both the LUB property and the NIP, and we will show that completeness of the real numbers implies the NIP. By Exercise 6.3.8, completeness also implies the LUB property.

One good reason an author might want to take completeness as an axiom is that it makes for a pretty neat way to fill in the spaces between the rational numbers with the irrationals and make the real number line into the smooth continuum we envision. It also is a natural way to arrive at a form of closure in the steps we take to build the real numbers from the set $\{0, 1\}$. In a nutshell, it goes like this.

In the way we described in Section 3.8, we begin with the set $\{0, 1\}$ and expand it to the whole numbers in order to achieve closure of addition. We then expand the whole numbers to the integers to achieve closure of additive inverses, and then we extend the integers to the rationals to have closure of multiplicative inverses for nonzero elements. Now envision the rationals as a set of points on the number line with all the irrationals missing, and consider the set of all possible sequences of rational numbers. A problem with the rationals is that many of the Cauchy sequences will not converge to a rational limit. There are a lot of unproved assumptions in the next example, but it illustrates the point.

Example 6.4.7 The terms

$$\begin{aligned}
a_0 &= 1\\
a_1 &= 1.4\\
a_2 &= 1.41\\
a_3 &= 1.414\\
&\vdots\\
a_{11} &= 1.41421356237\\
&\vdots
\end{aligned}$$

represent the start of a sequence that converges to $\sqrt{2}$. Every term in the sequence is rational, for it is a terminating decimal. The sequence is Cauchy, for, given $\epsilon > 0$, let N be a positive integer such that $10^{-N} < \epsilon$. Then if $m, n \geq N$, $|a_m - a_n| < \epsilon$ because a_m and a_n will agree through the Nth decimal place. ■

So the rational numbers are not a complete set, which means that they do not contain all possible limits of all Cauchy sequences. What are you going to do? If you want the rational numbers to be embedded in a larger set that is complete, you are going to have to throw in all possible limits of Cauchy sequences of rational numbers. With a little bit of work, you can fill in the holes in the number line with these irrational limits.

6.4.2 From Completeness to the NIP

If we can show that completeness of the real numbers implies the NIP, then we will have demonstrated the logical equivalence of the LUB property, the NIP, and completeness. Here are some reminders and a suggestion about how to prove the final exercise of this section.

In proving the NIP from the LUB property in Section 6.3, we supposed $\langle [a_n, b_n] \rangle$ is a nested sequence of closed intervals whose lengths tend to zero as $n \to \infty$. That the length of the intervals goes to zero did not by itself say anything about convergence of $\langle a_n \rangle$ or $\langle b_n \rangle$ separately. The intervals might get arbitrarily short, but if they are not nested, then nothing prevents them from waltzing forever up the real number line, so that neither $\langle a_n \rangle$ nor $\langle b_n \rangle$ converges. That the intervals are nested meant that $\langle a_n \rangle$ and $\langle b_n \rangle$ are bounded. We then applied monotonicity and the LUB property to the a_n and b_n, which put a real number limit into our hands. This limit was just the number we needed to arrive at the NIP.

This time, in proving the NIP from completeness, we still begin with the assumption that the intervals are nested and that their lengths tend to zero. The difference this time is that the axiom we apply is completeness. If we can show that $\langle a_n \rangle$ and $\langle b_n \rangle$ are Cauchy, then the assumption of completeness will give us a real number limit for each. This turns out to be fruitful, for the fact that $(b_n - a_n) \to 0$ makes $|b_n - a_n|$ eventually very small. To conclude that $\langle a_n \rangle$ is Cauchy, b_n in the expression $|b_n - a_n|$ can be replaced by a_m for $m \geq n$ to have $|a_m - a_n|$, which is even smaller than $|b_n - a_n|$. The rest of the work in proving the NIP would be identical to what you did in Section 6.3, where we observed $\lim a_n = \lim b_n$, and where you showed in Exercise 6.3.4 that $\cap[a_n, b_n] = \{L\}$, where L is the common limit of $\langle a_n \rangle$ and $\langle b_n \rangle$. So here is the only result we need to bridge the gap from completeness to the NIP.

EXERCISE 6.4.8 Suppose $\langle [a_n, b_n] \rangle$ is a sequence of nested intervals such that $(b_n - a_n) \to 0$ as $n \to \infty$. Then $\langle a_n \rangle$ is Cauchy. (A similar argument would show that $\langle b_n \rangle$ is Cauchy.)

Functions of a Real Variable

Functions whose domain is a subset of the real numbers are called *functions of a real variable*, and functions whose range is a subset of the real numbers are called *real-valued*. Some especially powerful and interesting results can be deduced if the domain of such a function is compact. Because this chapter deals exclusively with real-valued functions of a real variable, we will write $f : S \to \mathbb{R}$, where it is understood that S is a subset of the real numbers.

7.1 Bounded and Monotone Functions

First we discuss boundedness and monotonicity of a function $f : S \to \mathbb{R}$. With these and other ideas from Sections 7.2–7.4, we will point out many links to properties of sequences we discussed in Chapter 6.

7.1.1 Bounded Functions

The definitions of the boundedness terms for functions resemble those for sequences.

Definition 7.1.1 Suppose $f : S \to \mathbb{R}$ is a function and A is a subset of S. We say f is *bounded from above* on A provided there exists a real number M_1 such that $f(x) \le M_1$ for all $x \in A$. Similarly, we say f is *bounded from below* on A provided there exists a real number M_2 such that $f(x) \ge M_2$ for all $x \in A$. If there exists $M > 0$ such that $|f(x)| \le M$ for all $x \in A$, we say f is *bounded* on A. For each of these characteristics, if it applies on the entire domain, we say simply that f is bounded from above, bounded from below, or bounded.

Exercise 7.1.2 Use the logic of Definition 7.1.1 to state what the following terms mean.

(a) f is not bounded from above on A.

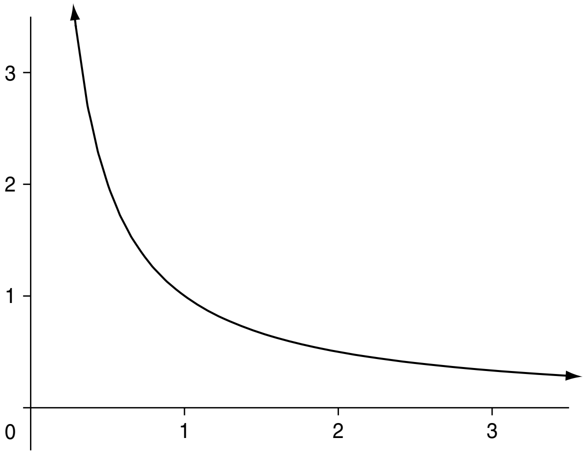

Figure 7.1 $f(x) = 1/x$.

(b) f is not bounded from below on A.

(c) f is not bounded on A.

Exercise 7.1.3 For $f(x) = 1/x$ on the interval $(0, \infty)$ show the following. (See Fig. 7.1.)

(a) f is bounded on $(1, \infty)$.

(b) f is bounded from below on $(0, 1)$.

(c) f is not bounded from above on $(0, 1)$.

7.1.2 Monotone Functions

When we defined a sequence to be monotone increasing, the fact that the terms of the sequence are lined up in an order made the definition easy to come by—if we pick two indices $m < n$, it must be that $a_m \leq a_n$. A similar definition works for functions.

Definition 7.1.4 Suppose $f : S \to \mathbb{R}$ is a function, A is a subset of S, and $a, b \in A$. Then f is said to be *increasing* (or *nondecreasing*) on A provided $f(a) \leq f(b)$ whenever $a < b$. Similarly, f is said to be *decreasing* (or *nonincreasing*) on A provided $f(a) \geq f(b)$ whenever $a < b$. If $f(a) < f(b)$ whenever $a < b$, f is said to be *strictly increasing* on A. If $f(a) > f(b)$ whenever $a < b$, f is said to be *strictly decreasing* on A. For each of these characteristics, if it applies on the entire domain, we say simply that f is increasing, decreasing, strictly increasing, or strictly decreasing, and we group functions of these types into the class we call *monotone* functions.

Exercise 7.1.5 Show $f(x) = 1/(x^2 + 1)$ is strictly decreasing on $[0, \infty)$.

In Exercise 3.5.7, you proved the following theorem, even though we did not have the language of monotonicity at the time.

Theorem 7.1.6 Let n be any positive integer. Then:

1. $f(x) = x^{2n-1}$ is strictly increasing on the set of real numbers.
2. $f(x) = x^{2n}$ is strictly decreasing on $(-\infty, 0]$ and strictly increasing on $[0, \infty)$.

In your pre-calculus work, you probably learned about the *horizontal line test*, an intuitive way of determining whether a function f is one-to-one. If every horizontal line in the xy-plane crosses the graph of f no more than once, then f is one-to-one. If f does not map onto the real numbers, we can crop the codomain of f down to Rng f so that $f : S \to \text{Rng } f$ is onto. Thus if $f : S \to \text{Rng } f$ is one-to-one, then $f^{-1} : \text{Rng } f \to S$ will exist. To find the expression for f^{-1}, you switched the roles of x and y in the expression for f, then solved for y. You also learned that the graph of f^{-1} could be sketched by reflecting the graph of f about the diagonal line $y = x$. This reflection is the visual result of swapping x and y. Think intuitively for a moment. If f is a strictly increasing function, then does it pass the horizontal line test? If so, then f^{-1} exists on Rng f. Given this, can you say anything about the monotonicity of f^{-1}?

Theorem 7.1.7 If a function is strictly monotone on A, then it is one-to-one on A.

Exercise 7.1.8 Prove Theorem 7.1.7 for the case of an increasing function. (A similar argument would work for a decreasing function.)

As an immediate consequence of Exercise 7.1.8 and Theorem 4.4.7, we have the following.

Corollary 7.1.9 If f is strictly monotone on $A \subseteq S$, then there exists an inverse function $f^{-1} : f(A) \to A$.

Letting $A = S$ in Corollary 7.1.9, we have that if $f : S \to \mathbb{R}$ is strictly monotone, then $f^{-1} : \text{Rng } f \to S$ exists. Furthermore, the following is true.

Theorem 7.1.10 If $f : S \to \mathbb{R}$ is strictly increasing (or strictly decreasing), then $f^{-1} : \text{Rng}(f) \to S$ is strictly increasing (or strictly decreasing).

Exercise 7.1.11 Prove the case of Theorem 7.1.10 where f is increasing. (A similar argument would work if f is decreasing.)

Theorems 7.1.6 and 7.1.10 imply that $\sqrt[2n]{x}$ is strictly increasing on $[0, \infty)$ and $\sqrt[2n-1]{x}$ is strictly increasing on $\mathbb{R}$. Also, with the next exercise, the function $f(x) = x^r$

is strictly increasing on $[0, \infty)$ for any rational number r. For we may write $r = m/n$ for positive integers m and n, so that $f(x) = \sqrt[n]{x^m}$.

Exercise 7.1.12 Suppose f and g are real-valued functions such that $g \circ f$ is defined on some set S. If f and g are both increasing, then so is $g \circ f$.

Exercise 7.1.13 State and prove other results analogous to Exercise 7.1.12 by varying the hypothesis conditions on f and g to include all combinations of increasing and decreasing.

7.2 Limits and Their Basic Properties

Suppose f is a function that is defined on some deleted neighborhood of a. Whether $f(a)$ exists is beside the point in the definition of limit, so we insist on no more than a deleted neighborhood. In a way somewhat similar to our definition of convergence of sequences as $n \to \infty$, we discuss $\lim_{x\to a} f(x)$, the limit of $f(x)$ as x approaches a. In this section, we will work our way into the definition of the limit of a function, and then look at some basic theorems involving limits that bear striking parallels to those theorems involving sequences.

7.2.1 Definition of Limit

Let's work our way into a definition of $\lim_{x\to a} f(x)$ gradually. Let $f(x)$ be defined in the following way:

$$f(x) = \begin{cases} 4x - 1, & \text{if } x \neq 2 \\ 50, & \text{if } x = 2 \end{cases} \tag{7.1}$$

Defining f in this way produces a very familiar straight line, but with a hole punched in the graph at $x = 2$ and the value of $f(2)$ artificially (but purposefully) defined to be a number way out of line with what you would expect it to be. (See Fig. 7.2.) Even though $f(2) = 50$, the function does not take on values anywhere near 50 if x is close to but different from 2. To the contrary, if x is in some small *deleted* neighborhood of 2, $f(x)$ appears to take on values close to 7. The next example asks the question that will lead us into the definition of limit.

Example 7.2.1 Let f be defined as in Eq. (7.1), and suppose $\epsilon > 0$ is given. For the ϵ-neighborhood of 7 on the y-axis in Figure 7.2, we want to find a radius $\delta > 0$ of a deleted neighborhood of 2 on the x-axis so that every x in the deleted neighborhood maps into $N_\epsilon(7)$ on the y-axis. That is, we want to find $\delta > 0$ so that $|f(x) - 7| < \epsilon$ whenever $0 < |x - a| < \delta$.

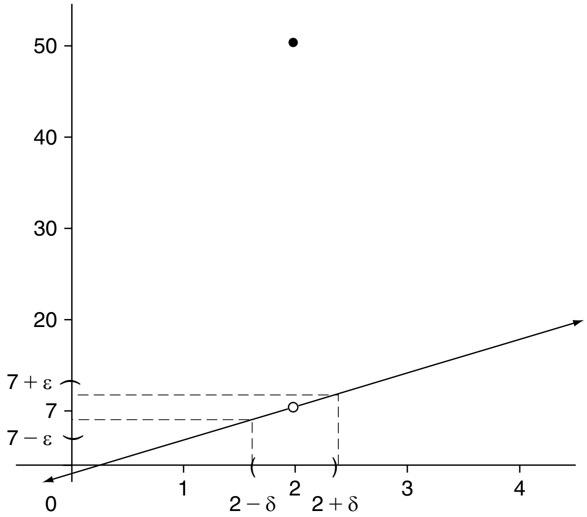

Figure 7.2 f defined from Eq. (7.1).

Solution Since the only values of x we care about are different from 2, we may use $f(x) = 4x - 1$ in the inequality $|f(x) - 7| < \epsilon$. Thus we have the following equivalent inequalities.

$$7 - \epsilon < 4x - 1 < 7 + \epsilon$$

$$2 - \frac{\epsilon}{4} < x < 2 + \frac{\epsilon}{4}$$

$$|x - 2| < \frac{\epsilon}{4}$$

By letting $\delta = \epsilon/4$ we guarantee that if $0 < |x - 2| < \delta$, then $|f(x) - 7| < \epsilon$. ■

Example 7.2.1 illustrates the following point, which is at the heart of the definition of $\lim_{x\to a} f(x)$. Given a "target" of radius ϵ around $y = 7$, we are able to find a deleted neighborhood around $x = 2$ such that all values of the function in the deleted neighborhood of $x = 2$ hit somewhere in the target neighborhood of $y = 7$. If you imagine a sequence of smaller and smaller ϵ-values, it is always possible to find sufficiently small deleted neighborhoods of $x = 2$ that map entirely within the smaller and smaller targets. If such a deleted δ-neighborhood can be found for all ϵ-neighborhoods of $y = 7$, then we say that $\lim_{x\to 2} f(x) = 7$. Here is the definition.

Definition 7.2.2 Suppose $f : S \to \mathbb{R}$ is a function defined on some deleted neighborhood of a, and L is a real number. Then we say $\lim_{x\to a} f(x) = L$ provided for all $\epsilon > 0$, there exists $\delta > 0$ such that, for all $x \in S$, $0 < |x - a| < \delta$ implies

$|f(x) - L| < \epsilon$. If $\lim_{x\to a} f(x) = L$, we say f *converges* to L as $x \to a$. Another way to write this is $f(x) \to L$ as $x \to a$.

To write Definition 7.2.2 symbolically, we would have

$$\lim_{x\to a} f(x) = L \Leftrightarrow (\forall\epsilon > 0)(\exists\delta > 0)(\forall x \in S)(0 < |x - a| < \delta \to |f(x) - L| < \epsilon) \tag{7.2}$$

Here is an example of a proof based on Definition 7.2.2. Some scratchwork similar to the inequalities in Example 7.2.1 has been omitted.

Example 7.2.3 Let g be defined by

$$g(x) = \frac{3x^2 + 17x + 20}{2x + 8} \tag{7.3}$$

Note that the natural domain of g is all real numbers except $x = -4$. Show $\lim_{x\to -4} g(x) = -7/2$.

Solution Let $\epsilon > 0$ be given, and let $\delta = 2\epsilon/3$. Now if $x \neq -4$, a factor of $(x + 4)$ may be canceled from the numerator and denominator of g. Thus if $0 < |x + 4| < \delta$, we have

$$\begin{aligned}\left|g(x) + \frac{7}{2}\right| &= \left|\frac{3x^2 + 17x + 20}{2x + 8} + \frac{7}{2}\right| \\ &= \left|\frac{(3x + 5)(x + 4)}{2(x + 4)} + \frac{7}{2}\right| \\ &= \left|\frac{3x + 5}{2} + \frac{7}{2}\right| = \left|\frac{3x + 12}{2}\right| \\ &= \left|\frac{3(x + 4)}{2}\right| = \frac{3}{2}\,|x + 4| < \frac{3\delta}{2} = \epsilon\end{aligned} \tag{7.4}$$

Thus $\lim_{x\to -4} g(x) = -7/2$. ■

In the same way that f from Example 7.2.1 seemed a little contrived because of the special effort we took to define $f(2) = 50$, g from Example 7.2.3 might seem contrived because nothing would seem to prevent us from canceling the $(x + 4)$ factor from numerator and denominator. Granted, g is merely the straight line $y = 3x + 5$ with a hole punched in the domain at $x = -4$ by the introduction of an $(x + 4)$ factor in the numerator and denominator. But f and g are good first examples to illustrate that the limit of a function such as $x \to a$ has nothing to do with the value or even the existence of the function at $x = a$. Whether the function exists at a and is the same as the limit as $x \to a$ is called *continuity*, and we'll look at that in Section 7.5. There are plenty of examples in which a hole in the domain occurs naturally, and we will look at one example later in this section.

7.2.2 Basic Theorems of Limits

There is a real similarity between the definition of convergence of a sequence as $n \to \infty$ and convergence of a function as $x \to a$. In the convergence of sequences, a given $\epsilon > 0$ must be associated with a threshold term beyond which all terms of the sequence fall in the ϵ-neighborhood of the limit. In the convergence of a function f, a given $\epsilon > 0$ must be associated with a δ-radius of a deleted neighborhood of a, within which all values of the function, except perhaps $f(a)$, must fall in the ϵ-neighborhood of the limit. The good news is that the similarity of these definitions makes for plenty of similar theorems, with similar proofs to match. So here is a barrage of theorems involving convergence of a function. As we present them, we will often link them back to sequence theorems and properties. We prove some of them here to illustrate the similarities, but you will prove most of them as exercises by mimicking your work from Section 6.2.

Theorem 7.2.4 For any real numbers a and c, the constant function $f(x) = c$ satisfies $\lim_{x\to a} f(x) = c$.

Proof. Let $\epsilon > 0$ be given, and let $\delta = 1$. Then if $0 < |x - a| < \delta$, $|f(x) - c| = |c - c| = 0 < \epsilon$. Thus $\lim_{x\to a} f(x) = c$. □

Exercise 7.2.5 For any real number a, $\lim_{x\to a} x = a$.

Theorem 7.2.6 If $f(x) \to L$ as $x \to a$, then there exists a deleted neighborhood of a where f is bounded.

Proof. Suppose $f(x) \to L$ as $x \to a$. Then there exists $\delta > 0$ such that $0 < |x - a| < \delta$ implies that $|f(x) - L| < 1$. Let $M = 1 + |L|$. Then $M > 0$ and if $0 < |x - a| < \delta$, we have

$$|f(x)| = |f(x) - L + L| \le |f(x) - L| + |L| < 1 + |L| = M$$

□

Exercise 7.2.7 If $f(x) \to 0$ as $x \to a$, and if g is bounded on a deleted neighborhood of a, then $f(x)g(x) \to 0$ as $x \to a$.

Exercise 7.2.8 If $f(x) \to L_1$ and $g(x) \to L_2$ as $x \to a$, then as $x \to a$:

(a) $f(x) + g(x) \to L_1 + L_2$

(b) $f(x)g(x) \to L_1 L_2$

Three results follow as immediate corollaries of Theorem 7.2.4 and Exercise 7.2.8.

Corollary 7.2.9 If $f(x) \to L$ as $x \to a$, and if c is any real number, then $cf(x) \to cL$ as $x \to a$.

Corollary 7.2.10 If $f(x) \to L_1$ and $g(x) \to L_2$ as $x \to a$, then $f(x) - g(x) \to L_1 - L_2$ as $x \to a$.

Corollary 7.2.11 If $P(x) = a_n x^n + a_{n-1}x^{n-1} + \cdots + a_1 x + a_0$, then for any real number a, $P(x) \to P(a)$ as $x \to a$.

We state the next result in a way slightly different from the corresponding result for sequences.

Exercise 7.2.12 Suppose $g(x) \to L$ as $x \to a$, where L is nonzero. Then there exists a deleted neighborhood of a where g can be bounded away from zero. In particular, if $L > 0$, there exists $M > 0$ and $\delta > 0$ such that $0 < |x - a| < \delta$ implies $g(x) > M$. If $L < 0$, there exists $M < 0$ and $\delta > 0$ such that $0 < |x - a| < \delta$ implies $g(x) < M$.

Exercise 7.2.13 Suppose $g(x) \to L$ as $x \to a$, where L is nonzero. Then $1/g(x) \to 1/L$ as $x \to a$.

Corollary 7.2.14 Suppose $f(x) \to L_1$ and $g(x) \to L_2 \neq 0$ as $x \to a$. Then $f(x)/g(x) \to L_1/L_2$ as $x \to a$.

Theorem 7.2.15 Suppose f is defined by

$$f(x) = \frac{P_1(x)}{P_2(x)} = \frac{a_r x^r + a_{r-1}x^{r-1} + \cdots + a_1 x + a_0}{b_s x^s + b_{s-1}x^{s-1} + \cdots + b_1 x + b_0}$$

Then, provided $P_2(a) \neq 0$, $f(x) \to P_1(a)/P_2(a)$ as $x \to a$.

Proof. Apply Corollaries 7.2.11 and 7.2.14, and the result follows. □

If $f(x) \to L$ and $g(x) \to L$ as $x \to a$, then $f(x) - g(x) \to 0$. Thus if $\epsilon > 0$, there exists $\delta > 0$ such that $0 < |x - a| < \delta$ implies $|f(x) - g(x)| < \epsilon/2$. This fact will come in handy in proving the following.

Exercise 7.2.16 [Sandwich Theorem] Suppose f, g, and h are functions with the property that $g(x) \leq h(x) \leq f(x)$ for all x in some deleted neighborhood of a. Suppose also that $f(x) \to L$ and $g(x) \to L$ as $x \to a$. Then $h(x) \to L$ as $x \to a$.

Earlier we promised an example of a function with a natural hole in the domain where we want to address the limit. The function is $f(x) = \sin x/x$ at

$a = 0$. Except for a passing glance at sine in Section 6.1, we have never mentioned any trigonometric, exponential, or logarithmic functions. There is a reason for this: Definitions of these functions that are rooted in the axioms of the real numbers are not possible to come by at this stage of our game. Only the algebraic functions arise from the theory we have developed thus far, and they can make for some pretty sticky proofs by themselves. Strict definitions of these nonalgebraic functions, called *transcendental functions*, come later. However, if we kick back for a while and give ourselves the freedom to talk about sine, cosine, and tangent in the familiar language of the unit circle, then we can apply Exercise 7.2.16 to show

$$\lim_{x\to 0} \frac{\sin x}{x} = 1 \tag{7.5}$$

In trigonometry, we take a real x-number line and "wrap" it around the unit circle in the uv-plane with $x = 0$ placed at $(u, v) = (1, 0)$ and the positive half of the x-axis wrapped counterclockwise. (See Fig. 7.3.) For a real number x, we define $\cos x$ and $\sin x$ to be the u and v coordinates, respectively, of the point where x falls in the uv-plane. If $0 < x < \pi/2$, Figure 7.3 shows how we may view x as an arc length, and $\sin x$, $\cos x$, and $\tan x$ as the lengths of segments in Quadrant I. Similar triangles in the figure will convince you that the length labeled $\tan x$ is correct. If $0 < x < \pi/2$ as is suggested in the sketch, the geometry of the unit circle reveals

$$\sin x \le x \le \tan x \tag{7.6}$$

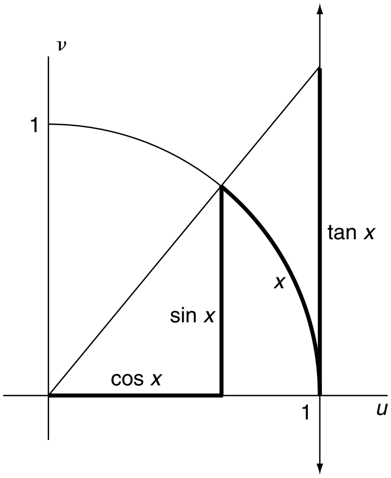

Figure 7.3 Basic trigonometric functions.

If $-\pi/2 < x < 0$ so that x is in Quadrant IV, then $\sin x$ and $\tan x$ are also negative, and we have

$$\tan x \leq x \leq \sin x \tag{7.7}$$

If we take both parts of the inequality in (7.6) separately and solve each for $\sin x/x$, we can put them back together into the single inequality

$$\cos x \leq \frac{\sin x}{x} \leq 1 \tag{7.8}$$

If we do the same thing for (7.7), remembering that $x < 0$ and $\sin x < 0$, we also arrive at (7.8), so that (7.8) holds for all $0 < |x| < \pi/2$. Let's accept from the geometry of the unit circle and the definition of $\cos x$ that $\cos x \to 1$ as $x \to 0$. Then, by Exercise 7.2.16, Eq. (7.5) holds.

EXERCISE 7.2.17 Assume that $-1 \leq \cos x \leq 1$ for all x, and show that $\lim_{x\to 0} x \cos x = 0$.

Let's remind ourselves of some of the theory of Cauchy sequences in order to motivate our last theorem from this section. An immediate result of convergence of a sequence is that it is Cauchy (Exercise 6.4.2). If the terms of a sequence get close to some limit, then they must get close to each other. Similarly, we can make the following claim, which says that if $f(x)$ converges as $x \to a$, then the values of f must not vary much from each other in small deleted neighborhoods of a.

EXERCISE 7.2.18 Suppose $f(x) \to L$ as $x \to a$. Then for every $\epsilon > 0$, there exists $\delta > 0$ such that $x_1, x_2 \in DN_\delta(a)$ implies $|f(x_1) - f(x_2)| < \epsilon$.

If the conclusion statement in Exercise 7.2.18 does not hold, then there exists $\epsilon > 0$ such that for all $\delta > 0$, there will exist $x_1, x_2 \in DN_\delta(a)$ with $|f(x_1) - f(x_2)| \geq \epsilon$. That is, there is some $\epsilon > 0$ such that every deleted δ-neighborhood of a contains two points whose functional values are at least ϵ apart. If a function f has this property, then $\lim_{x\to a} f(x)$ does not exist.

Example 7.2.19 Let f be defined in the following way:

$$f(x) = \begin{cases} 1, & \text{if } x \text{ is rational} \\ 0, & \text{if } x \text{ is irrational} \end{cases} \tag{7.9}$$

If a is any real number, then $\lim_{x\to a} f(x)$ does not exist. For we may let $\epsilon = 1/2$ and choose any $\delta > 0$. Then by Exercises 5.2.4 and 5.2.5, the deleted δ-neighborhood of a contains a rational number x_1 and an irrational number x_2, and $|f(x_1) - f(x_2)| = 1 > \epsilon$. This function is sometimes called the *salt and pepper* function because its values of 0 and 1 are sprinkled up and down the domain like grains of salt and pepper. No one knows for sure which is the salt and which is the pepper. ■

Exercise 7.2.20 The *signum* function, $\operatorname{sgn} x$, is defined in the following way:

$$\operatorname{sgn} x = \begin{cases} -1, & \text{if } x < 0 \\ 0, & \text{if } x = 0 \\ 1, & \text{if } x > 0 \end{cases} \tag{7.10}$$

Show $\lim_{x\to 0} \operatorname{sgn} x$ does not exist.

Exercise 7.2.21 Use familiar values of $\sin x$ to show that $\lim_{x\to 0} \sin(1/x)$ does not exist.

7.3 More on Limits

7.3.1 One-Sided Limits

In defining $\lim_{x\to a} f(x) = L$, we insisted f be defined on both sides of a and all points to the nearby left and right of a map into $N_\epsilon(L)$. If f is only defined on one side of a, or perhaps if values of f to the immediate left of a behave differently from those to the immediate right (as in $\operatorname{sgn} x$), we can discuss the limit of $f(x)$ as x approaches a from the left or from the right separately. Instead of using entire deleted neighborhoods of a, we use only the left or right half of them.

Definition 7.3.1 Suppose f is defined on the interval (a, b). Then we say $\lim_{x\to a^+} f(x) = L$ (read "as x approaches a from the right") if for all $\epsilon > 0$, there exists $\delta > 0$ such that $a < x < a + \delta$ implies $|f(x) - L| < \epsilon$. We call L the *right-hand limit* of f at a. Similarly, if f is defined on the interval (c, a), we say $\lim_{x\to a^-} f(x) = L$ (x approaches a from the left) if for all $\epsilon > 0$, there exists $\delta > 0$ such that $a - \delta < x < a$ implies $|f(x) - L| < \epsilon$. We call L the *left-hand limit* of f at a.

Exercise 7.3.2 Given a function f, $\lim_{x\to a} f(x) = L$ if and only if

$$\lim_{x\to a^+} f(x) = \lim_{x\to a^-} f(x) = L \tag{7.11}$$

Example 7.3.3 Let f be defined in the following way:

$$f(x) = \begin{cases} x^2 - 1, & \text{if } x < 2 \\ \frac{1}{x-1}, & \text{if } x > 2 \end{cases} \tag{7.12}$$

By Theorem 7.2.15, $\lim_{x\to 2}(x^2-1)=3$ and $\lim_{x\to 2} 1/(x-1)=1$ (in the two-sided sense). Thus, by the $\Rightarrow$ direction of Exercise 7.3.2, $\lim_{x\to 2^-}(x^2-1)=3$ and $\lim_{x\to 2^+} 1/(x-1)=1$. But then $\lim_{x\to 2^-} f(x)=3$ and $\lim_{x\to 2^+} f(x)=1$, so that $\lim_{x\to 2} f(x)$ fails to exist by the $\Leftarrow$ direction of Exercise 7.3.2. ■

Exercise 7.3.4 Let f be defined in the following way:

$$f(x)=\begin{cases} 2\operatorname{sgn} x, & \text{if } x\le 1 \\ x^3+1, & \text{if } x>1 \end{cases} \tag{7.13}$$

Evaluate the following, with reference to applicable results.

(a) $\lim_{x\to 0^-} f(x)$

(b) $\lim_{x\to 0^+} f(x)$

(c) $\lim_{x\to 0} f(x)$

(d) $\lim_{x\to 1^-} f(x)$

(e) $\lim_{x\to 1^+} f(x)$

(f) $\lim_{x\to 1} f(x)$

Exercise 7.3.5 Show $\lim_{x\to 0^+}\sqrt{x}=0$.

7.3.2 Sequential Limits

Suppose f is a function defined on a deleted neighborhood of a, where $f(x)\to L$ as $x\to a$. Let $\langle a_n\rangle$ be any sequence of real numbers that converges to a, but where every a_n is different from a. Now consider the sequence $\langle f(a_n)\rangle$. (See Fig. 7.4.) It seems intuitively clear that $f(a_n)\to L$ as $n\to\infty$, and this is a fairly straightforward fact to demonstrate.

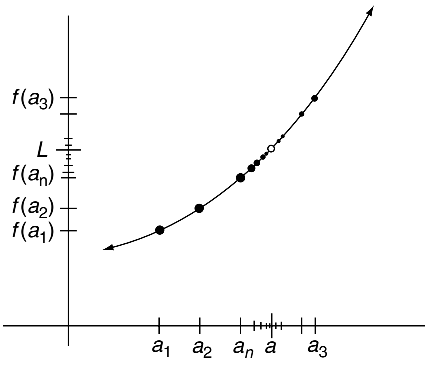

Figure 7.4 Convergence of a function in terms of sequences of domain elements.

What might be surprising is that the converse of this result is also true. If every sequence of real numbers converging to a (where all the a_n are different from a) produces a sequence of functional values that converges to L, then $f(x) \to L$ as $x \to a$. One way to prove this is by contrapositive. Suppose $f(x)$ does not converge to L as $x \to a$, claim the ϵ that this guarantees, and then use a sequence of progressively smaller and smaller δ-values to create a sequence $\langle a_n \rangle$ where all a_n are different from a and that converges to a, but where $f(a_n)$ does not converge to L.

EXERCISE 7.3.6 Suppose $f : S \to \mathbb{R}$ is a function. Then $\lim_{x \to a} f(x) = L$ if and only if every sequence $\langle a_n \rangle$ such that $a_n \to a$ and $a_n \neq a$, for all n also has the property that $f(a_n) \to L$.

The logical equivalence of our ϵ-δ definition of limit and the sequential limit characteristic in Exercise 7.3.6 means that it is possible to define the statement $\lim_{x \to a} f(x) = L$ using either an ϵ-δ definition or a sequential limit definition. Some authors prefer to define function limits by defining $\lim_{x \to a} f(x) = L$ provided every sequence $\langle a_n \rangle$ where every $a_n \neq a$ and $a_n \to a$ as $n \to \infty$ implies that $f(a_n) \to L$. Our ϵ-δ feature would then be a theorem. They then construct proofs of theorems such as those from Section 7.2 from the theory of sequences we discussed in Section 6.2. If you already have the theory of sequences under your belt, function limit proofs become very easy. For example, let's return to Exercise 7.2.8(a): If $f(x) \to L_1$ and $g(x) \to L_2$ as $x \to a$, then $f(x) + g(x) \to L_1 + L_2$. A proof using a sequential limit definition and exploiting the theorems about sequences would go something like this.

Proof of Exercise 7.2.8(a). Suppose $f(x) \to L_1$ and $g(x) \to L_2$ as $x \to a$. Let $\langle a_n \rangle$ be a sequence such that $a_n \to a$ and $a_n \neq a$ for all n. Then since $f(x) \to L_1$ as $x \to a$, we have that $f(a_n) \to L_1$ as $n \to \infty$. Similarly, since $g(x) \to L_2$ as $x \to a$, then $g(a_n) \to L_2$ as $n \to \infty$. By Exercise 6.2.12, $f(a_n) + g(a_n) \to L_1 + L_2$. Since $\langle a_n \rangle$ was chosen arbitrarily, we have $f(x) + g(x) \to L_1 + L_2$ as $x \to a$. □

Perhaps you feel a little cheated at this point. After all, if only we had proved Exercise 7.3.6 at the beginning of Section 7.2, then all our proofs about limits would have been one-liners. There is some truth to that. However, ϵ-δ proofs pervade mathematics, and the fact that your first ϵ-δ proofs could be done easily by merely mimicking work from Section 6.2 was probably a humane way for you to be introduced to them.

EXERCISE 7.3.7 Use Exercise 7.3.6 to show that $\lim_{x \to 0} \sin(1/x)$ does not exist by demonstrating an appropriate sequence.

7.4 Limits Involving Infinity

Now we want to extend the concepts and language of limits to include the symbols $\pm\infty$, even though infinity is not a real number. We have used the symbol for infinity before in our work with sequences when we defined $\lim_{n\to\infty} a_n$. However, our use of the symbol was merely a formal one, because convergence was defined solely in terms of the indexing set of positive integers. In this section, we want to take the expression $\lim_{x\to a} f(x) = L$ and replace either a or L (or both) with one of the symbols $\pm\infty$. Up until this point, the fact that a and L are real numbers allowed us to discuss ϵ-neighborhoods of L and δ-neighborhoods of a. With our current definition of neighborhood, it makes no sense to talk about a *neighborhood of infinity*, unless of course we decide to give this term meaning. We will do precisely this. In fact, as we go, we will concoct extended, new meanings for old language and revamp some of our visual imagery to show that there just might be some way to understand infinity in a way that we can almost treat it as a number.

The language we will use in replacing a or L with $\pm\infty$ will go something like this. In discussing $\lim_{x\to\pm\infty} = L$, we call these limits *at* positive or negative infinity. When we discuss $\lim_{x\to a} f(x) = \pm\infty$, we call these limits *of* positive or negative infinity. Graphically, a limit at infinity corresponds to a *horizontal asymptote* in the graph of f, and a limit of infinity corresponds to a *vertical asymptote*. There are buckets and buckets of ways to combine and specialize the ideas we will discuss here. With both $+\infty$ and $-\infty$, with either a or L or both being replaced by these symbols, with two-sided and one-sided limits, we can define a whole bunch of new terms. By hitting a few, you will catch on to what the others ought to be, so we will not be exhaustive.

7.4.1 Limits at Infinity

Let's begin by replacing a with $+\infty$, because it almost exactly replicates the theory of sequences. Since a sequence is just a real-valued function whose domain is the positive integers, our way of graphing a sequence as we did in Figure 6.1 makes convergence as $n \to \infty$ easy to visualize as a horizontal asymptote of discrete points in the plane. If we imagine filling in values of the function to other real numbers so that the domain is some interval $(a, +\infty)$, then the following definition seems to be a natural adaptation of Definition 6.2.1.

Definition 7.4.1 Suppose f is defined on some interval $(a, +\infty)$. Then we say $\lim_{x\to+\infty} f(x) = L$ provided for all $\epsilon > 0$, there exists a real number M such that $x > M$ implies $|f(x) - L| < \epsilon$.

Notice the similarity between Definitions 7.4.1 and 6.2.1. Given $\epsilon > 0$, there exists a threshold point in the domain beyond which all values of the function fall in the ϵ-neighborhood of L. This definition gives rise to a whole slew of theorems involving limits of functions at $+\infty$, where the proofs are identical to those for sequences.

Theorem 7.4.2 $\lim_{x\to+\infty} c = c$.

Theorem 7.4.3 $\lim_{x\to+\infty} 1/x = 0$.

Theorem 7.4.4 Suppose $f(x) \to L_1$ and $g(x) \to L_2$ as $x \to +\infty$. Then the following hold as $x \to +\infty$.

1. $f(x) + g(x) \to L_1 + L_2$
2. $f(x)g(x) \to L_1L_2$
3. $f(x)/g(x) \to L_1/L_2$ (if $L_2 \neq 0$)

On and on the theorems go that exactly parallel our previous work. The limit of polynomial over polynomial, and even a sandwich theorem, seem strangely translucent.

EXERCISE 7.4.5 Create a definition for $\lim_{x\to-\infty} f(x) = L$ for a function defined on an appropriate interval.

Now let's extend our language of neighborhood to include infinity. When we say $f(x) \to L$ as $x \to a$, we mean that any neighborhood of L has a corresponding deleted neighborhood of a that maps into it. Is there a way to use the same language for a limit at infinity? Could we say that any neighborhood of L has a corresponding neighborhood of $+\infty$ that maps into it? We can of course, if we define a *neighborhood of* $+\infty$ to be an interval of the form $(M, +\infty)$. Similarly, a neighborhood of $-\infty$ could be defined as an interval of the form $(-\infty, M)$.

We are catching a glimpse of the *extended real numbers* and are developing an imagery of two phantom points $\pm\infty$ somewhere way off the left and right ends of the number line. This is the standard way of creating the extended real numbers, which are denoted as either $\overline{\mathbb{R}}$ or $[-\infty, +\infty]$.

An apparent difference between neighborhoods of a real number and neighborhoods of $\pm\infty$ is that the latter are not two-sided. However, if we use the single symbol ∞ instead of both $\pm\infty$, then we can create a nice way of visualizing the extended real numbers, where $+\infty$ and $-\infty$ are merged into one point. Here is one way we might do that, which is similar to the extension of the xy-plane to include the *point at infinity* in complex analysis.

Imagine a real number line with a circle sitting on top of it as in Figure 7.5. We map a given real number a to a corresponding point (x, y) on the circle with the help of the diagonal line in the figure. This geometric way of mapping each

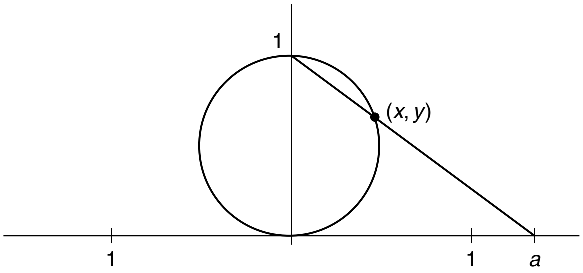

Figure 7.5 Mapping the extended real numbers onto a circle.

real number to a point on the circle suggests a one-to-one function from the real numbers *onto* all points of the circle except the north pole $(0, 1)$. If you want an explicit formula for the coordinates of the point (x, y) in terms of the value of a, you can get it easily enough with the help of similar triangles and the equation for the circle, but it is not necessary to understand the principle.

EXERCISE 7.4.6 Use similar triangles and the equation of the circle from Figure 7.5 to determine the xy-coordinates of the image of a real number a under the bijection that sends the real numbers to all points on the circle except $(0, 1)$.

With this new imagery, the real numbers are no longer conceived as an infinite line but as a circle, except that one point of the circle has not been associated with a real number in this mapping. We call the point at $(0, 1)$ the point at infinity.

The extension of the real numbers to $\overline{\mathbb{R}}$ by introducing positive and negative infinity as strange "end points" of the number line is the more common conception of the extended real numbers. The point is that $\pm\infty$ are nothing more than formal symbols thrown in, along with a bunch of algebraic rules about how you are supposed to use them. We have therefore arrived at the place where $\lim_{x\to a} f(x) = L$ has meaning for all real numbers L and all $-\infty \le a \le +\infty$.

7.4.2 Limits of Infinity

Now let's replace L with $\pm\infty$ in the expression $\lim_{x\to a} f(x) = L$.

Definition 7.4.7 Suppose f is defined on some deleted neighborhood of a. We say $\lim_{x\to a} f(x) = +\infty$ provided for all $M > 0$, there exists $\delta > 0$ such that $0 < |x - a| < \delta$ implies $f(x) > M$.

Definition 7.4.7 is a two-sided definition, so the graph of a function f for which $\lim_{x\to a} f(x) = +\infty$ will have a vertical asymptote at a, both sides of which head upward.

Example 7.4.8 Show $\lim_{x\to 0} 1/x^2 = +\infty$.

Solution Pick $M > 0$. Let $\delta = 1/\sqrt{M}$. Then if $0 < |x| < \delta$, it follows that

$$f(x) = \frac{1}{x^2} = \frac{1}{|x|^2} > \frac{1}{\delta^2} = M \tag{7.14}$$

■

EXERCISE 7.4.9 In the spirit of Definitions 7.4.1 and 7.4.7, create definitions for the following statements.

(a) $\lim_{x\to a} f(x) = -\infty$

(b) $\lim_{x\to a^+} f(x) = +\infty$

(c) $\lim_{x\to a^-} f(x) = +\infty$

(d) $\lim_{x\to +\infty} f(x) = +\infty$

(e) $\lim_{x\to +\infty} f(x) = -\infty$

(f) $\lim_{x\to -\infty} f(x) = +\infty$

(g) $\lim_{x\to -\infty} f(x) = -\infty$

EXERCISE 7.4.10 Given a function f, $\lim_{x\to a} f(x) = +\infty$ if and only if

$$\lim_{x\to a^+} f(x) = \lim_{x\to a^-} f(x) = +\infty \tag{7.15}$$

What other kinds of theorems can we expect for limits of infinity? Can we prove something that resembles Theorem 7.4.4? The answer is yes, but the theorems will look different because the limits are not necessarily real numbers. Consequently, we have to begin with Definition 7.4.7 and do some of the work from scratch. The parts of the next exercise are ordered so that some of the later ones follow quickly from the earlier ones. This should make your work a little more efficient. These results and those in Exercise 7.4.12 could just as easily be stated and proved in terms of one-sided limits.

EXERCISE 7.4.11 Suppose $f(x) \to L$, $g(x) \to +\infty$, and $h(x) \to +\infty$ as $x \to a$. Then as $x \to a$,

(a) $-g(x) \to -\infty$

(b) $1/g(x) \to 0$

(c) $f(x) + g(x) \to +\infty$

(d) $f(x)g(x) \to +\infty$ if $L > 0$

(e) $f(x)g(x) \to -\infty$ if $L < 0$

(f) No conclusion can be drawn about $f(x)g(x)$ if $L = 0$.

(g) $f(x)/g(x) \to 0$

(h) $g(x) + h(x) \to +\infty$

(i) $g(x)h(x) \to +\infty$

(j) No conclusion can be drawn about $g(x) - h(x)$ or $g(x)/h(x)$.

If $f(x) \to 0$ as $x \to a$, then we can say something about $1/f(x)$ under certain circumstances.

Exercise 7.4.12 Suppose there exists a deleted neighborhood of a such that $f(x) > 0$ for all x in the deleted neighborhood. Suppose also that $f(x) \to 0$ as $x \to a$. Then $1/f(x) \to +\infty$. Similarly, if $f(x) < 0$ for all x in the deleted neighborhood and $f(x) \to 0$, then $1/f(x) \to -\infty$.

In the same way that we use the notation $x \to a^+$ and $x \to a^-$ to mean x approaches a from the right and left, respectively, we can create a shorthand notation for functions that behave as in Exercise 7.4.12. We write $f(x) \to 0^+$ to mean that f approaches zero and is positive on a deleted neighborhood of a, and we say f approaches zero through positive values. Similarly we write $f(x) \to 0^-$ to mean that f approaches zero and is negative on a deleted neighborhood of a, and we say f approaches zero through negative values.

Exercise 7.4.13 Let $f(x) = (x^2 - 4)/(x - 1)$ for all real numbers $x \neq 1$. Find with verification $\lim_{x\to 1^-} f(x)$ and $\lim_{x\to 1^+} f(x)$.

7.5 Continuity

The word *continuous* may already be a part of your mathematical vocabulary. Perhaps your calculus class delved into continuity enough to provide an ϵ-δ definition. More than likely, your notions of continuity are probably best summarized as a belief that a continuous function can be sketched without lifting your pencil off the paper. The graph is one clean, easily drawable piece. Even though such a view can be helpful in your understanding of some characteristics of continuity, it is far from true that all continuous functions are so easily drawable. The bizarre examples of undrawable continuous functions will come later in your study of mathematics. For now, we define the terms and study the basic results. We begin with continuity at a single point in the domain, and then we talk about continuity on a subset of the domain. In the same way that we talk about left-hand and right-hand limits, we will then talk about left continuity and right continuity.

7.5.1 Continuity at a Point

If $\lim_{x\to a} f(x)$ exists, it describes the behavior of f near a but not at a. It is possible that $\lim_{x\to a} f(x) = L$, while $f(a)$ might either fail to exist or exist and be different from L. If a is in the interior of the domain of f, if $\lim_{x\to a} f(x)$ exists, and if this limit is the value of $f(a)$, we give this phenomenon a name.

Definition 7.5.1 Suppose $f : S \to \mathbb{R}$ is a function and $a \in \text{Int}(S)$. We say that f is *continuous at a* provided

$$\lim_{x\to a} f(x) = f(a) \tag{7.16}$$

If f is not continuous at a, we say that f is *discontinuous* there or that f has a *discontinuity* at a.

Let's reword Definition 7.5.1 in the language of ϵ, δ, and neighborhoods. Continuity is a small step logically from limit, for all we do is delete the word *deleted* when we discuss neighborhoods of a. Definition 7.5.1 can be reworded to say that f is continuous at a provided

$$(\forall \epsilon > 0)(\exists \delta > 0)(\forall x \in S)(|x - a| < \delta \rightarrow |f(x) - f(a)| < \epsilon) \tag{7.17}$$

Compare (7.17) to the symbolic form of the definition of limit in (7.2). The only difference is that $|x - a| < \delta$ replaces $0 < |x - a| < \delta$. And since $|x - a| < \delta$ is *weaker* than $0 < |x - a| < \delta$, continuity is therefore *stronger* than the existence of a limit.

Because of the short step from limit to continuity, our work with limits in Section 7.2 takes us a long way in the theory of continuous functions. If we go back to the results from Section 7.2, replacing L with $f(a)$ and converting deleted neighborhoods to neighborhoods, we arrive immediately at the following theorems.

Theorem 7.5.2 The constant function $f(x) = c$ is continuous at every real number.

Theorem 7.5.3 The identity function is continuous at every real number.

Theorem 7.5.4 If f is continuous at a point, then there exists a neighborhood of the point where f is bounded.

Theorem 7.5.5 If f and g are both continuous at a point, then so are $f + g$, fg, and $f - g$.

Theorem 7.5.6 Suppose g is continuous at a and $g(a) \neq 0$. Then there exists a neighborhood of a where g can be bounded away from zero. In particular, if $g(a) > 0$, there exists $M > 0$ and $\delta > 0$ such that $|x - a| < \delta$ implies $g(x) > M$. If $g(a) < 0$, there exists $M < 0$ and $\delta > 0$ such that $|x - a| < \delta$ implies $g(x) < M$.

Theorem 7.5.7 If g is continuous at a and $g(a) \neq 0$, then $1/g$ is continuous at a.

Theorem 7.5.8 If f and g are both continuous at a and if $g(a) \neq 0$, then f/g is continuous at a.

Corollary 7.5.9 A polynomial function is continuous at every real number. Every rational function $f(x) = P_1(x)/P_2(x)$, where P_1 and P_2 are polynomials, is continuous at every real number a for which $P_2(a) \neq 0$.

Theorem 7.5.10 A function f is continuous at a if and only if every sequence $\langle a_n \rangle$ such that $a_n \rightarrow a$ satisfies $f(a_n) \rightarrow f(a)$.

Notice that the hypothesis condition of Theorem 7.5.10 does not require $a_n \neq a$ as Exercise 7.3.6 does. For if any of the $a_n = a$, then $f(a_n) = f(a)$, so that the sequence $\langle f(a_n) \rangle$ is defined for all n.

Here is an important result we could not address with limits alone. A composition function $g \circ f$ is continuous at a point if f and g are continuous at the right points. To prove this, an arbitrarily chosen ϵ-neighborhood of $g[f(a)]$ produces a δ_1-neighborhood of $f(a)$, and this δ_1-neighborhood of $f(a)$ produces a δ_2-neighborhood of a.

EXERCISE 7.5.11 Suppose f is continuous at a and g is continuous at $f(a)$. Then $g \circ f$ is continuous at a.

Another way to write what Exercise 7.5.11 says is

$$\lim_{x\to a} g(f(x)) = g(\lim_{x\to a} f(x)) = g(f(\lim_{x\to a} x) = g(f(a)) \tag{7.18}$$

With a little more mathematical machinery than we have up to now, Eq. (7.18) comes in handy in certain manipulations of limits. If we assume for the moment that all the functions below are continuous at the points involved, Eq. (7.18) allows us to write something like

$$\begin{aligned} \lim_{x\to 3} \sqrt{1+e^{-x^2}} &= \sqrt{\lim_{x\to 3}(1+e^{-x^2})} \\ &= \sqrt{1+\lim_{x\to 3} e^{-x^2}} \\ &= \sqrt{1+e^{\lim_{x\to 3}(-x^2)}} \\ &= \sqrt{1+e^{-9}} \end{aligned} \tag{7.19}$$

We could not have a theorem in Section 7.2 that said something like "If $f(x) \to L_1$ as $x \to a$ and $g(x) \to L_2$ as $x \to L_1$, then $(g \circ f)(x) \to L_2$ as $x \to a$" because a possible hole in the domain of g at $f(a)$ might cause troublesome gaps in the domain of $g \circ f$ around a.

Example 7.5.12 Let $f(x) = x\cos(1/x)$ and $g(x) = \sin x/x$, and let $a = 0$. Now $\cos(1/x)$ is bounded, so that $f(x) \to 0$ as $x \to 0$. Also, by the presence of $\cos(1/x)$, $f(x) = 0$ for infinitely many values of x in every neighborhood of zero. Thus $(g \circ f)(x)$ fails to exist at infinitely many points in every neighborhood of zero because $g(0)$ is not defined. ■

If f is discontinuous at a point it could be for one or more of three basic reasons. If Eq. (7.16) is not satisfied it might be that

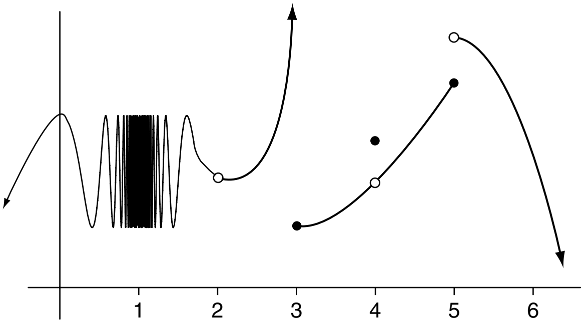

Figure 7.6 Some examples of discontinuities.

(D1) $\lim_{x\to a} f(x)$ does not exist;

(D2) $f(a)$ does not exist; or

(D3) $\lim_{x\to a} f(x)$ and $f(a)$ both exist, but are not equal.

Figure 7.6 illustrates some of these possibilities at $a = 1, 2, 3, 4, 5$. At $a = 1$, f behaves much like $\sin(1/x)$ near $x = 0$. (See Exercise 7.2.21.)

Exercise 7.5.13 For the function sketched in Figure 7.6, state which of the characteristics D1–D3 describe the discontinuities at $x = 1, 2, 3, 4, 5$.

The discontinuities at 2 and 4 in Figure 7.6 are called *removable* discontinuities because it is possible to define or redefine the value of f to make it continuous there. The discontinuity at 5 is called a *jump discontinuity*.

Example 7.5.14 The function $f(x) = \sin x/x$ has a removable discontinuity at zero. Thus the function

$$g(x) = \begin{cases} \dfrac{\sin x}{x}, & \text{if } x \neq 0 \\ 1, & \text{if } x = 0 \end{cases} \tag{7.20}$$

is continuous at zero because $\lim_{x\to 0} f(x) = 1 = f(0)$. ■

Example 7.5.15 The salt and pepper function (Example 7.2.19) is discontinuous at every real number, for $\lim_{x\to a} f(x)$ fails to exist for every real number a. ■

The next exercise presents a well-known function that is related to the salt and pepper function. It provides an interesting first glimpse into functions in which continuity and discontinuity can coexist in functions whose behavior might not seem so odd at first glance.

Exercise 7.5.16 Define $f : (0, 1) \to \mathbb{R}$ in the following way.

$$f(x) = \begin{cases} \dfrac{1}{q}, & \text{if } x = p/q \text{ is rational } (p \text{ and } q \text{ relatively prime}) \\ 0, & \text{if } x \text{ is irrational} \end{cases} \tag{7.21}$$

Then f is continuous at every irrational and discontinuous at every rational.[1]

7.5.2 Continuity on a Set

The definition of continuity on a subset of the domain is pretty straightforward if the subset is open, for then we are guaranteed that every point is contained in a neighborhood entirely within the domain of the function.

Definition 7.5.17 If $f : S \to \mathbb{R}$ is a function and A is an open subset of S, we say f is *continuous on* A provided it is continuous at every point in A. Logically, we may write this as

$$(\forall a \in A)(\forall \epsilon > 0)(\exists \delta > 0)(\forall x \in A)(|x - a| < \delta \to |f(x) - f(a)| < \epsilon) \tag{7.22}$$

Notice that the only difference between statements (7.17) and (7.22) is that the latter begins with $(\forall a \in A)$. To construct a proof that f is continuous on a set A by going back to ϵ and δ, you would begin by picking both $a \in A$ and $\epsilon > 0$. We want to do precisely this in an example now. We want to show that $f(x) = 1/x$ is continuous on $(0, +\infty)$. Yes, it can be said that continuity of $f(x) = 1/x$ on $(0, +\infty)$ follows from Theorem 7.5.9. But we are trying to illustrate that finding a value of δ will depend not only on the value of ϵ, but on the arbitrarily chosen point in the domain at which we are working.

To begin the construction of a proof that $f(x) = 1/x$ is continuous on $(0, +\infty)$, pick $a > 0$ and $\epsilon > 0$. As usual in an ϵ-δ proof, we must find $\delta > 0$ so that the inequality

$$|x - a| < \delta \tag{7.23}$$

will be at least as strong an inequality as

$$\left| \frac{1}{x} - \frac{1}{a} \right| < \epsilon \tag{7.24}$$

So we work backwards from (7.24), trying to transform it into something of the form (7.23), making sure that the steps we take do not produce inequalities that get any weaker. One way to rewrite (7.24) is

[1] For any positive integer n, there are only finitely many rational numbers with denominator $q \leq n$.

$$\frac{|a-x|}{|xa|} = \frac{|x-a|}{|x|\,|a|} < \epsilon \tag{7.25}$$

Now a is a given positive number, and provided we choose δ to be sufficiently small, x will be sufficiently close to a and therefore be positive also. Furthermore, to make inequality (7.25) true, we must bound $|x-a|$ from above, and we must bound $|x|$ from below by some number, so that we can replace $|x|$ with this number in the inequality. We must take both of these considerations into account to find δ, and here is one way to do that.

First, we can insist that the neighborhood around a be sufficiently small so that $x > a/2$. This is equivalent to $a - x < a/2$, so we insist that $\delta < a/2$. Second, assuming $x > a/2$, we would have $1/x > 2/a$, to have

$$\frac{|x-a|}{a^2/2} < \epsilon \tag{7.26}$$

Multiplying this through by $a^2/2$ suggests that $|x-a|$ needs to be smaller than $\epsilon a^2/2$. Thus if we let $\delta = \min\{a/2, \epsilon a^2/2\}$, then we are sure that δ is positive and that inequality (7.25) will be satisfied. A proof that $1/x$ is continuous on $(0, +\infty)$ would then go something like this.

Proof. Let a and ϵ be given positive numbers. Let $\delta = \min\{a/2, \epsilon a^2/2\}$, which is clearly positive. Then if $|x-a| < \delta$, we note first that $x > a - \delta \geq a/2$ because $\delta \leq a/2$. Furthermore, since $\delta \leq \epsilon a^2/2$,

$$\left|\frac{1}{x} - \frac{1}{a}\right| = \frac{|a-x|}{|xa|} = \frac{|x-a|}{|x|\,a} < \frac{|x-a|}{a^2/2} < \epsilon \tag{7.27}$$

Thus $1/x$ is continuous at a. □

Notice a very important fact. The value of δ we proposed does depend on both ϵ and a. In fact, the closer a is to zero, the smaller our value of δ will be. On the basis of our work, it appears there is no single value of δ that would depend only on ϵ and work for all positive a. It might occur to you, however, that someone else might have taken a different approach and stumbled across a value of δ that was not tied to the value of a but depended only on ϵ.

The question is an important one and will be answered when we study uniform continuity in Section 7.7. Suffice it to say for now that no such δ as a function of ϵ alone exists for $f(x) = 1/x$ on $(0, +\infty)$. Intuitively, it might seem plausible when we consider the vertical asymptote of $1/x$ at zero. In Figure 7.7, imagine setting an ϵ-tolerance around a point $1/a$ on the y-axis. If a is very close to zero, the graph of f is very steep at the point $(a, 1/a)$ and varies a lot even in a small neighborhood of a on the x-axis. Granted, for any a there is a δ-neighborhood of a that maps into $N_\epsilon[f(a)]$. But if you imagine smaller and smaller values of a, then the asymptote of the graph necessitates smaller and smaller values of δ. No single δ-value will work for all a because of this asymptote.

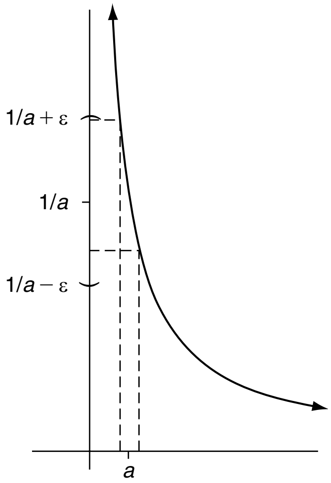

Figure 7.7 $f(x) = 1/x$.

The scratchwork you will go through in the next exercise is not nearly as involved as it has been for our example here, but you will learn a bit more about using the inequality $f(a) - \epsilon < f(x) < f(a) + \epsilon$ and working backwards to find an inequality of the form $a - \delta < x < a + \delta$ that is at least as strong. A peculiar thing for $\sqrt{x}$ on $(0, \infty)$, however, is that it actually is possible to find a δ that is independent of the value of a. Though it is true that the graph of $\sqrt{x}$ gets very steep as x approaches zero from the right, it does not behave asymptotically there. Why such a δ can be found will become clear in Section 7.7.

EXERCISE 7.5.18 Use an ϵ-δ proof to show that $f(x) = \sqrt{x}$ is continuous at any $a > 0$.

7.5.3 One-Sided Continuity

Even if $\lim_{x\to a} f(x)$ does not exist, there might still be a hope that either $\lim_{x\to a^-} f(x)$ or $\lim_{x\to a^+} f(x)$ exists. Similarly, though f might not be continuous at a, we might have continuity from one side or the other.

Definition 7.5.19 A function f is said to be *left continuous* at a provided

$$\lim_{x\to a^-} f(x) = f(a) \tag{7.28}$$

Similarly, f is said to be *right continuous* at a provided

$$\lim_{x\to a^+} f(x) = f(a) \tag{7.29}$$

An immediate consequence of Exercise 7.3.2 is the following.

Theorem 7.5.20 A function f is continuous at a if and only if it is both left and right continuous at a.

In defining continuity on a set, Definition 7.5.17 stipulated that the set must be open. This allowed for a definition of continuity on (a, b), but not on $[a, b]$. One-sided continuity now allows us to define continuity on $[a, b]$ in such a way that we do not concern ourselves with how f behaves or whether it even exists outside of $[a, b]$.

Definition 7.5.21 A function f is said to be continuous on $[a, b]$ provided the following hold:

1. f is continuous on (a, b),
2. f is right continuous at a,
3. f is left continuous at b.

Definition 7.5.21 can be naturally adapted to apply to the following example by omitting stipulation 3.

Example 7.5.22 $f(x) = \sqrt{x}$ is continuous on $[0, +\infty)$. From Exercise 7.5.18, f is continuous at all positive real numbers, and by Exercise 7.3.5, it is right continuous at 0.

7.6 Implications of Continuity

Continuity has many implications, and in this section we look at three of them. The second of these results is actually logically equivalent to continuity.

7.6.1 The Intermediate Value Theorem

The imagery that we can sketch a continuous function without picking up the pencil makes this first result seem plausible. It says that a continuous function cannot be negative at one point and positive somewhere else without crossing the x-axis somewhere between the two points. One case of this is stated in the next exercise. To prove it, imagine standing on the x-axis at a and looking up the x-axis. There is a natural subset of the real numbers that is both non-empty and bounded from above by b to which the LUB property can apply.

Exercise 7.6.1 Suppose $a < b$, $f(a) < 0 < f(b)$, and f is continuous on $[a, b]$. Then there exists $c \in (a, b)$ such that $f(c) = 0$.[2]

[2] Let $A = \{x \in [a, b] : f(x) < 0\}$. What can you say about A?

If g is a continuous function such that $g(a) > 0 > g(b)$ for some $a < b$, then we can apply Exercise 7.6.1 to $-g$ to produce some $c \in (a, b)$ such that $g(c) = 0$.

Exercise 7.6.1 provides most of the machinery needed to prove a theorem whose name you might remember from calculus, the Intermediate Value Theorem (IVT). We will supply the proof here to illustrate a convenience, where a specific result like Exercise 7.6.1 can generalize very easily.

Theorem 7.6.2 (Intermediate Value Theorem). Suppose f is continuous on $[a, b]$, and suppose $f(a) < f(b)$. Let y_0 be any real number satisfying $f(a) < y_0 < f(b)$. Then there exists $c \in (a, b)$ such that $f(c) = y_0$.

In the IVT, f does not necessarily change from negative to positive as x runs from a to b. However, if we create a new function g by dropping or raising f to make $g(a) < 0 < g(b)$, then we can apply Exercise 7.6.1 to g and see what it says about f.

Proof. Define $g(x) = f(x) - y_0$. Since f and the constant function y_0 are both continuous on $[a, b]$, Theorem 7.5.5 and our work with one-sided continuity imply that g is continuous on $[a, b]$. Furthermore, $g(a) = f(a) - y_0 < 0$ and $g(b) = f(b) - y_0 > 0$. By Exercise 7.6.1, there exists $c \in (a, b)$ such that $g(c) = 0$. Thus $f(c) = y_0$. □

EXERCISE 7.6.3 Prove the *fixed point theorem*: If $f : [a, b] \to [a, b]$ is a continuous function, then there exists $c \in [a, b]$ such that $f(c) = c$.[3]

EXERCISE 7.6.4 The *fundamental theorem of algebra* says every polynomial function whose degree is odd has a real root. That is, if $P(x) = a_{2n+1}x^{2n+1} + a_{2n}x^{2n} + \cdots + a_1x + a_0$, then there exists a real number c such that $P(c) = 0$. In this exercise you prove the fundamental theorem of algebra, first for the case $a_{2n+1} > 0$ and then for the case $a_{2n+1} < 0$.

(a) Use your definitions from Exercise 7.4.9 to prove

$$\lim_{x\to+\infty} x = +\infty \quad \text{and} \quad \lim_{x\to-\infty} x = -\infty \tag{7.30}$$

(b) With part (a) in hand, a result like Exercise 7.4.11 would be demonstrable for $x \to \pm\infty$ by paralleling its proof. Assuming this result, show that $P(x)$ as defined above and with $a_{2n+1} > 0$ satisfies

$$\lim_{x\to-\infty} P(x) = -\infty \quad \text{and} \quad \lim_{x\to+\infty} P(x) = +\infty \tag{7.31}$$

[3] Consider the function $g(x) = f(x) - x$. If either $g(a)$ or $g(b)$ is zero, you're done. Otherwise, apply the IVT.

(c) Use your result from part (b) to prove the fundamental theorem of algebra for the case $a_{2n+1} > 0$.

(d) Prove the fundamental theorem of algebra for the case $a_{2n+1} < 0$ by applying part (c) to $-P(x)$.

Theorem 7.1.7 and Corollary 7.1.9 demonstrated that strict monotonicity of a function on a set implies that it is one-to-one (hence invertible on the image of the set). The converse of Corollary 7.1.9 is clearly not true, as is illustrated by $f(x) = 1/x$ on the nonzero real numbers, which is one-to-one but not monotone. However, if f is one-to-one and continuous on a set, then it must be strictly monotone. The IVT will come in handy in the proof of the next result.

EXERCISE 7.6.5 If f is continuous and one-to-one on $[a, b]$, then f is strictly monotone on $[a, b]$.[4,5]

The IVT can also help us prove that a continuous, invertible function has a continuous inverse. Specifically, if f is continuous and invertible on $[a, b]$, then it is continuous and one-to-one. By Exercise 7.6.5, f is strictly monotone on $[a, b]$ also. If f is strictly increasing, then $f(a) < f(x) < f(b)$ for all $x \in (a, b)$. By the IVT, every $y \in [f(a), f(b)]$ has a pre-image in $[a, b]$, so that $f([a, b]) = [f(a), f(b)]$. By similar reasoning, if f is strictly decreasing, $f([a, b]) = [f(b), f(a)]$. In either case, $f([a, b]) = [c, d]$ for some $c < d$. To show f^{-1} is continuous on $[c, d]$, you will have to verify all three criteria in Definition 7.5.21.

EXERCISE 7.6.6 Suppose f is continuous and invertible on $[a, b]$. Then $f^{-1} : f([a, b]) \to [a, b]$ is continuous on $f([a, b])$.

7.6.2 Continuity and Open Sets

The logical equivalence of the ϵ-δ form of continuity and the sequential limit form conveyed by Theorem 7.5.10 are useful, not only because they give us freedom to exchange one property for another, but also because they suggest an alternative way to define continuity that might be preferred in the creation of some mathematical structures. Exercise 7.6.7 states another feature that is logically equivalent to continuity, this time in terms of the pre-images of open subsets of the codomain. In the theorems that follow, we are going to assume that the functions involved are defined on the entire set of real numbers. This will keep the proofs relatively simple and get the point across, though similar theorems can be addressed on restricted domains.

[4] See Exercise 1.2.18(i).

[5] Follow your nose. For f not monotone and $f(a) < f(b)$, spend some time showing there exist $c_1 < c_2 < c_3$ such that $f(c_1) < f(c_2) > f(c_3)$ or $f(c_1) > f(c_2) < f(c_3)$.

When you prove Exercise 7.6.7, you can write a very elegant proof if you will use a slightly different form of the definition of continuity than that in expression (7.17). The following statements are all equivalent:

$$\begin{aligned} |x-a| < \delta &\to |f(x) - f(a)| < \epsilon \\ x \in N_\delta(a) &\to f(x) \in N_\epsilon[f(a)] \\ f[N_\delta(a)] &\subseteq N_\epsilon[f(a)] \qquad (7.32) \end{aligned}$$

If you will use (7.32) as you apply continuity, you will find that Exercises 4.3.6, 4.3.10, and 4.3.13 make for some pretty nifty manipulations of the sets involved in the proof.

EXERCISE 7.6.7 A function $f : \mathbb{R} \to \mathbb{R}$ is continuous if and only if the pre-image of every open set is open.

With Exercise 7.6.7, the following should drop right into your lap.

EXERCISE 7.6.8 A function $f : \mathbb{R} \to \mathbb{R}$ is continuous if and only if the pre-image of every closed set is closed.[6]

Let's digress just a moment to comment on the significance of Exercise 7.6.7. First, the logical equivalence of the two statements in Exercise 7.6.7 suggests yet another place one might begin in defining continuity of a function. If one were to begin by defining a function to be continuous provided the pre-image of every open set is open, then our ϵ-δ definition and the sequential limit property of Exercise 7.3.6 would become theorems in the analysis of real numbers. If it seems a little unnatural to begin with such an open set pre-image definition, consider the following, which is a glimpse into topology.

At the beginning of Chapter 5, we said that a defining characteristic of analysis is that elements of a set have either a measure of size (norm) or of distance between them (metric). In the real numbers, the measure most commonly used is absolute value, so that $|x|$ is the size of x and $|a - b|$ is the distance between a and b. If a set is endowed with a metric, then a neighborhood of a is defined as all points within a certain distance of a. Then a definition of open set like Definition 5.3.2 becomes meaningful because openness is defined in terms of neighborhoods.

In topology there is no such notion of size or distance either defined or assumed on the set. Instead, we begin with the idea of open set in what might seem a peculiar way. Given a set S, we declare some subsets to be open just because we say so. If you want to declare every subset of S to be open, fine. However, that might not prove to be especially interesting. To prevent your unbridled freedom to define open sets from degenerating into mathematically useless anarchy, your collection of open sets needs to have some of the same properties of open sets we derived as theorems for the real numbers. Specifically, if we lump all the subsets of S that

[6] See Exercise 4.3.11(c).

we declare to be open into the family $\mathcal{O}$, then in forming a topology we insist on the following.

(T1) Both $\emptyset$ and S are sets in the family $\mathcal{O}$.

(T2) If $\mathcal{F} \subseteq \mathcal{O}$ is a collection of open sets, then $\bigcup_{A\in\mathcal{F}} A$ is open. That is, $\mathcal{O}$ is closed under union.

(T3) If $\{A_k : 1 \le k \le n\}$ is a finite collection of open sets, then $\bigcap_{k=1}^{n} A_k$ is open. That is, $\mathcal{O}$ is closed under finite intersection.

Closed sets are then defined to be those whose complement is open, and you are on your way to building a structure that will have some parallels to those we have studied for the real numbers. But the structure will be more abstract and austere because the definition of open set does not probe as deeply into some assumed structure of the set.

Exercise 7.6.7 will also help you prove the following.

EXERCISE 7.6.9 The continuous image of a compact set is compact. That is, if $f : \mathbb{R} \to \mathbb{R}$ is continuous on a compact set S, then $f(S)$ is compact.

In Exercise 5.4.7 , you showed that if L is the LUB of a set S and $L \notin S$, then L is a cluster point of S. Since a compact set is bounded, it has an LUB, and since it is also closed, it contains all its cluster points. Thus a compact set contains its LUB. Similarly, a compact set also contains its GLB. If $f : \mathbb{R} \to \mathbb{R}$ is continuous and S is compact, then $f(S)$ is compact by Exercise 7.6.9. Consequently, $f(S)$ contains its LUB and GLB. Write $M = \max[f(S)]$ and $m = \min[f(S)]$. Then there exist $x_1, x_2 \in S$ such that $f(x_1) = m$ and $f(x_2) = M$. We have just proved the following theorem.

Theorem 7.6.10 (Extreme Value Theorem (EVT)). If $f : \mathbb{R} \to \mathbb{R}$ is continuous on a compact set S, then f attains a maximum and minimum value on S.

EXERCISE 7.6.11 Give two examples to illustrate that compactness is necessary for the EVT to apply to a continuous function, one example where S is bounded but not closed, and one where S is closed but not bounded.

7.7 Uniform Continuity

In Section 5.2 we used a metaphor involving diseases, medication, and panaceas to distinguish between every question having an answer and the existence of a single answer that applies to every question. From Definition 7.5.17, continuity of a function f on a set A can be written logically as

$$(\forall a \in A)(\forall \epsilon > 0)(\exists \delta > 0)(\forall x \in A)(|x - a| < \delta \to |f(x) - f(a)| < \epsilon) \qquad (7.33)$$

In general, the value of δ will depend on both ϵ and a. In this section, we want to define a form of continuity in which δ can be found that depends only on ϵ and works for all $a \in A$. We will look at some examples to illustrate the point, and then we will study two important theorems.

7.7.1 Definition and Examples

Let's take the definition of continuity in (7.33) and leapfrog the first component piece $(\forall a \in A)$ two jumps to the right.

$$(\forall \epsilon > 0)(\exists \delta > 0)(\forall a \in A)(\forall x \in A)(|x - a| < \delta \to |f(x) - f(a)| < \epsilon) \tag{7.34}$$

Moving $(\forall a \in A)$ to the right of $(\forall \epsilon > 0)$ does not change anything logically from (7.33). However, moving $(\forall a \in A)$ to the right of $(\exists \delta > 0)$ makes a big difference. If you were writing a proof of continuity from (7.33), you would begin by picking $a \in A$ and $\epsilon > 0$. Then, with both a and ϵ in hand, you would go hunting for $\delta > 0$ with the required properties. But with the phrase $(\forall a \in A)$ repositioned as in (7.34), things are different. If you were writing a proof of a theorem where (7.34) was involved, you would begin by picking only $\epsilon > 0$. Then, without knowing any particular value of $a \in A$, you would have to find $\delta > 0$ from ϵ alone, and this δ would have to serve for all $a \in A$. So (7.34) suggests that the value of δ can be found after having specified only ϵ, and this δ-value will work for all $a, x \in A$ satisfying $|x - a| < \delta$. Because (7.33) specifies a first, we think of a as being the center of a δ-neighborhood and x as an arbitrary point in that neighborhood. There is no reason to call one point a and the other one x, as if one were fixed before the other. The point is that there exists $\delta > 0$ such that if any two points are within δ of each other, then their functional values are within ϵ of each other. For clarity and convenience, we change the symbols slightly in the following definition.

Definition 7.7.1 A function $f : S \to \mathbb{R}$ is said to be *uniformly continuous* on $A \subseteq S$ provided for all $\epsilon > 0$, there exists $\delta > 0$ such that, for all $x, y \in A$, $|x - y| < \delta$ implies $|f(x) - f(y)| < \epsilon$. Logically, we may write this as

$$(\forall \epsilon > 0)(\exists \delta > 0)(\forall x, y \in A)(|x - y| < \delta \to |f(x) - f(y)| < \epsilon) \tag{7.35}$$

To show a function f is uniformly continuous on A, we work backwards from $|f(x) - f(y)| < \epsilon$. Watch what happens in the next example.

Example 7.7.2 Show that $f(x) = x^2/(x + 1)$ is uniformly continuous on $[0, \infty)$.

Solution Before we write a demonstration, we need to do some scratchwork. If you play with the expression $|f(x) - f(y)| < \epsilon$ and try to factor $|x - y|$ out of it, you can arrive at the following.

$$|f(x) - f(y)| = \left|\frac{x^2}{x+1} - \frac{y^2}{y+1}\right| = |x - y|\left|\frac{xy + x + y}{(x+1)(y+1)}\right| \tag{7.36}$$

Next, let's split up the fraction in the right-hand side of Eq. (7.36) and apply the triangle inequality. Also, notice that $x, y > 0$ implies that $1/(x+1), 1/(y+1)$, $x/(x+1)$, and $y/(y+1)$ are all less than 1. So we have

$$\begin{aligned}\left|\frac{xy+x+y}{(x+1)(y+1)}\right| &\le \left|\frac{xy}{(x+1)(y+1)}\right| + \left|\frac{x}{(x+1)(y+1)}\right| + \left|\frac{y}{(x+1)(y+1)}\right| \\ &= \left|\frac{x}{(x+1)}\right|\left|\frac{y}{(y+1)}\right| + \left|\frac{x}{(x+1)}\right|\left|\frac{1}{(y+1)}\right| \\ &\quad + \left|\frac{y}{(y+1)}\right|\left|\frac{1}{(x+1)}\right| \le 3\end{aligned} \tag{7.37}$$

Having arrived at $|f(x) - f(y)| \le 3\,|x-y|$, we are ready to write a proof.

Let $\epsilon > 0$ be given, and let $\delta = \epsilon/3$. Then for any $x, y \ge 0$ such that $|x-y| < \delta$, we have

$$\begin{aligned}|f(x) - f(y)| &= \left|\frac{x^2}{x+1} - \frac{y^2}{y+1}\right| = |x-y|\left|\frac{xy+x+y}{(x+1)(y+1)}\right| \\ &\le |x-y|\left(\left|\frac{x}{(x+1)}\right|\left|\frac{y}{(y+1)}\right|\right. \\ &\quad \left. + \left|\frac{x}{(x+1)}\right|\left|\frac{1}{(y+1)}\right| + \left|\frac{y}{(y+1)}\right|\left|\frac{1}{(x+1)}\right|\right) \\ &\le 3\,|x-y| < 3\delta = \epsilon\end{aligned} \tag{7.38}$$

■

The scratchwork of Example 7.7.2 suggests the following more general result.

EXERCISE 7.7.3 If there exists $m > 0$ such that $|f(x) - f(y)| \le m\,|x-y|$ for all $x, y \in A$, then f is uniformly continuous on A.

If $y \ne x$, the hypothesis condition of Exercise 7.7.3 is equivalent to

$$\left|\frac{f(x) - f(y)}{x-y}\right| \le m \tag{7.39}$$

which means that f has a bound on the slopes of lines through any two points on its graph. Loosely speaking, if the steepness of f (as measured by slopes of secant lines) is bounded, then f is uniformly continuous. The converse of Exercise 7.7.3 is not true. Shortly, we will point out a function f that is uniformly continuous on a set A, but for which inequality (7.39) does not hold across A for any $m > 0$.

Naturally, we want to include a demonstration that a continuous function need not be uniformly continuous on a set.

EXERCISE 7.7.4 What does it mean for f not to be uniformly continuous on A?

Example 7.7.5 Show that $f(x) = 1/x$ is not uniformly continuous on $(0, 1)$.

Solution First, we do some scratchwork. The vertical asymptote at $x = 0$ will provide us with x and y values that are very close together but whose functional values can differ as much as a strategically chosen ϵ-value. We have to play with the inequality

$$\left|\frac{1}{x} - \frac{1}{y}\right| = \frac{|y - x|}{|x|\,|y|} \geq \epsilon \tag{7.40}$$

and find some $\epsilon > 0$ so that, regardless of $\delta > 0$, we can find two points x and y that are within δ of each other and satisfy inequality (7.40). No obvious ϵ-value jumps out at us, so let's try $\epsilon = 1$ and see if we can proceed.

Inequality (7.40) itself suggests a way to find x and y. Whatever we decide to let x be, we can let $y = x + \delta/2$ so that $|y - x| = \delta/2 < \delta$. Then the trick is to let x be sufficiently close to zero so that $|x|\,|y|$ is small enough to make inequality (7.40) true. Furthermore, since smaller values of δ represent our primary obstacle, we may assume δ is smaller than any convenient positive number, if such an assumption appears helpful. If we let $x = \delta/2$ and $y = \delta$, inequality (7.40) falls right into place, as long as $\delta < 1$. Here is our demonstration.

Let $\epsilon = 1$, and pick any $\delta > 0$. We may assume that $\delta < 1$. Let $x = \delta/2$ and $y = \delta$. Then $|x - y| < \delta$, and

$$|f(x) - f(y)| = \frac{|x - y|}{|x|\,|y|} = \frac{\delta/2}{\delta^2/2} = \frac{1}{\delta} > 1 = \epsilon \tag{7.41}$$

Thus f is not uniformly continuous on $(0, 1)$. ■

In Example 7.7.5, the fact that $(0, 1)$ is not closed allows for $1/x$ to have a vertical asymptote at one end point. Even though $1/x$ is continuous throughout $(0, 1)$, the asymptote is where we look to disprove uniform continuity. In the next exercise, you will need to look among sufficiently large numbers where you will find x and y within δ of each other whose functional values differ by at least ϵ.

EXERCISE 7.7.6 Show that $f(x) = x^2$ is not uniformly continuous on $[0, +\infty)$.

7.7.2 Uniform Continuity and Compact Sets

Perhaps the most beloved theorem dealing with uniform continuity is the following.

Theorem 7.7.7 If $f : S \to \mathbb{R}$ is a continuous function and S is compact, then f is uniformly continuous on S.

There is something about the fact that every open cover of S is reducible to a finite subcover that allows us to liberate the value of δ from any specific points in the domain and determine it from ϵ alone. We will supply the proof here because it requires a few sneaky shrinkings of neighborhoods. If you want to try to prove it on your own, here is a thumbnail sketch of how to proceed.

As usual, we pick $\epsilon > 0$. Then since f is continuous at every point in S, we can cover S with a slew of neighborhoods, one centered at each $a \in S$, whose radius is half the δ-value that guarantees $f[N_\delta(a)] \subseteq N_{\epsilon/2}[f(a)]$. Note that each value of δ will depend on a. Since every a is in its own δ-neighborhood, the set of all these neighborhoods covers S. Compactness of S then allows us to reduce this cover to a finite subcover. These finitely many neighborhoods supply us with a single δ-value: the minimum of the finitely many $\delta/2$-values from the subcover. We can then show that if $|x - y| < \delta$, then $|f(x) - f(y)| < \epsilon$. See if you can fill in the details. If not, here is the whole proof.

Proof. Let $\epsilon > 0$ be given. Since f is continuous at every $a \in S$, then for any particular a, there exists $\delta(a)$ (the notation illustrating that δ is a function of a), such that $|x - a| < \delta(a)$ implies $|f(x) - f(a)| < \epsilon/2$. Cover S with the set $\mathcal{C} = \{N_{\delta(a)/2}(a) : a \in S\}$. Since every a is in its own neighborhood of radius $\delta(a)/2$, $\mathcal{C}$ does in fact cover S. Since S is compact, C has a finite subcover $\mathcal{C}_1 = \{N_{\delta(a_k)/2}(a_k) : 1 \leq k \leq n\}$. Let $\delta = \min\{\delta(a_k)/2 : 1 \leq k \leq n\}$, and pick $x, y \in S$ such that $|x - y| < \delta$. Now since $\mathcal{C}_1$ covers S, there exists some k such that $x \in N_{\delta(a_k)/2}$. Furthermore, $y \in N_{\delta(a_k)}$ because

$$|y - a_k| \leq |y - x| + |x - a_k| < \delta + \frac{\delta(a_k)}{2} \leq \frac{\delta(a_k)}{2} + \frac{\delta(a_k)}{2} = \delta(a_k) \tag{7.42}$$

Thus

$$|f(x) - f(y)| \leq |f(x) - f(a_k)| + |f(a_k) - f(y)| < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon \tag{7.43}$$

so that f is uniformly continuous on S. □

In Exercise 7.5.18 you showed that $f(x) = \sqrt{x}$ is continuous on $(0, +\infty)$. Since $\lim_{x\to 0^+} \sqrt{x} = 0 = \sqrt{0}$, f is continuous on $[0, 1]$, and hence is uniformly continuous there. However, the graph of f gets very steep as $x \to 0^+$, so that (7.39) is not satisfied.

Exercise 7.7.8 Show that inequality (7.39) is not satisfied by $f(x) = \sqrt{x}$ on $[0, 1]$ for any $m > 0$.

PART III

Basic Principles of Algebra

This page intentionally left blank

Groups

In its simplest terms, algebra is the study of sets on which binary operations provide the defining internal structure for the set. For example, we may construct a set and define a form of addition or multiplication, then look at the structure of the set and the relationships between its elements that result from these operations.

Of particular interest in algebra is the study of mappings between sets S_1 and S_2 where the structure of the binary operation on S_1 is preserved among the images of the elements in S_2. More concretely, if S_1 has binary operation $*$, if S_2 has binary operation $\cdot$, and if f is a function from S_1 to S_2, we address the question of whether $f(a * b) = f(a) \cdot f(b)$ for all $a, b \in S_1$. Whether f is one-to-one or onto also leads to some interesting results.

In this chapter, we begin our study of some of the most basic concepts of algebraic structures by starting with groups. Some of the theorems are actually restatements of results we have already seen for the real numbers. Now, however, the context is broader, more abstract, so the air might seem a little thinner at first. Instead of proving algebraic theorems for which we have the real numbers specifically in mind, we prove theorems based on assumptions that certainly apply to real numbers, but of which the real numbers are only a specific example.

8.1 Introduction to Groups

8.1.1 Basic Characteristics of Algebraic Structures

An algebraic structure begins with a non-empty set and builds its internal structure in several stages. First there must be some notion of equality either defined or understood on the set, which naturally must be an equivalence relation. One of the assumptions we stated in Chapter 0 was that equality in the real numbers satisfies properties E1–E3, though we did not really probe into what is behind this assumed equality. Our work in Section 3.8 was a great example of how equality can be created in a set. Assuming that integer equality is an equivalence relation, we built the rationals using the integers as building blocks, where we defined rational equality in terms of integer equality. Then we showed that this definition is an equivalence relation.

Once the set is constructed and equality is defined, we can define one or more binary operations. By definition, a binary operation is well defined and closed. Let S be a set we have constructed, let $\equiv$ represent equality on S, and let $*$ be a binary operation defined on S. To say that the binary operation is well defined is to say that for all $a, b, c, d \in S$, if $a \equiv b$ and $c \equiv d$, then $a * c \equiv b * d$. Note how this requirement appeals to equality on S. From this point forward, we will almost always use $=$ to represent equality as it is defined on S.

Example 8.1.1 By assumption, addition and multiplication are well defined and closed on the real numbers, so that both are binary operations. Also by assumption, the positive real numbers are closed under addition and multiplication. ■

Example 8.1.2 By the remarks after Definition 3.8.6, addition is closed on the rationals. By Exercise 3.8.7, addition is well defined. Similarly, rational multiplication is closed, and by Exercise 3.8.10 it is well defined. ■

EXERCISE 8.1.3 Determine with explanation whether each of the following defines a binary operation on the given set.

(a) Addition on the negative real numbers.

(b) Multiplication on the negative real numbers.

(c) Addition on the irrational numbers.

(d) Multiplication on the irrational numbers.

(e) For a non-empty set U, the operation of union on the power set of U.

(f) For a non-empty set U, the operation of intersection on the power set of U.

(g) Greatest common divisor on the positive integers.

EXERCISE 8.1.4 For a given non-empty set A, let S be the set of all bijections from A to itself. Is composition a binary operation on S? Explain.

By our work in Section 3.11, a binary operation on S can be defined as a function $f : S \times S \to S$. The formal notation that would be used to represent addition as such a function would be $+ : S \times S \to S$, writing $+(a, b)$ to mean $a + b$. The fact that $+$ as a function has property F1 is precisely what it means for the binary operation of addition to be closed. That $+$ has property F2 means addition is well defined as an operation. Thus the fact that $+ : S \times S \to S$ is a function is equivalent to addition being a binary operation.

All the binary operations we will address in this chapter will have the associative property. We might have to verify associativity, however, if the context we are working in is new and we have created a binary operation from scratch. Some binary operations are not associative, but they are indeed rare.

Exercise 8.1.5 Give an example of a binary operation that is not associative.[1]

Many, but not all, of the binary operations we will consider will be commutative. Some very interesting results of algebra derive from binary operations that are not commutative. Be careful in your work. Unless commutativity is explicitly given or proved, you might be tempted to reverse the order of elements without any permission to do so.

Exercise 8.1.6 Show that the binary operation in Exercise 8.1.4 is not commutative.

The existence of an identity element for a binary operation is rather context specific. In the next definition we insist that an identity element must commute with every element of the set, regardless of whether the binary operation is commutative.

Definition 8.1.7 Suppose S is endowed with binary operation $*$. Then $e \in S$ is called an *identity* for the operation $*$ if $a * e = e * a = a$ for all $a \in S$.

If there is an identity element for a given binary operation, then it might be that some or all elements have an inverse.

Definition 8.1.8 Suppose S has binary operation $*$, for which there is an identity element e. If for a given $a \in S$ there exists $b \in S$ such that $a * b = b * a = e$, then we say that b is an *inverse* of a, and we write it as a^{-1}.

For convenience we sometimes list the features of an algebraic structure as an ordered list. For example, if S is a set with binary operation $*$, identity e, and with the feature that every $a \in S$ has an inverse a^{-1}, then we might write such a structure as $(S, *, e, ^{-1})$.

Of the basic algebraic properties we assumed on the real numbers (A2–A14), the only one we have not mentioned yet is the distributive property (A14). If a set is endowed with two binary operations, it might be that they are linked in their behavior by the distributive property. We will address algebraic structures with two binary operations in Chapter 9.

If S is a small finite set, it might be convenient to describe the binary operation in a *Cayley table*. Since a binary operation might not be commutative, it is important to read $a * b$ from a Cayley table by going down the left column to find a and across the top to find b.

[1] See Section 4.8.

Example 8.1.9 Consider $S = \{0, 1, 2, 3, 4, 5\}$ and let $\oplus$ be described as in Table 8.1.

$\oplus$	0	1	2	3	4	5
0	0	1	2	3	4	5
1	1	2	3	4	5	0
2	2	3	4	5	0	1
3	3	4	5	0	1	2
4	4	5	0	1	2	3
5	5	0	1	2	3	4

(8.1)

That $\oplus$ is well defined is immediate, for there is a unique value in each position in the table. Closure of $\oplus$ is also obvious, since every entry in the table is an element of S. ■

Exercise 8.1.10 Is $\oplus$ is commutative? How can you tell? Is there an identity element? Does every element have an inverse?

One way to verify associativity of $\oplus$ would be to do all possible calculations of the form $(a \oplus b) \oplus c$ and $a \oplus (b \oplus c)$ to check if they are equal. In Section 8.3, we will construct this algebraic structure formally and prove associativity as a theorem.

Example 8.1.11 Let S be the set in Example 8.1.9 and define the operation $\otimes$ as in Table 8.2.

$\otimes$	0	1	2	3	4	5
0	0	0	0	0	0	0
1	0	1	2	3	4	5
2	0	2	4	0	2	4
3	0	3	0	3	0	3
4	0	4	2	0	4	2
5	0	5	4	3	2	1

(8.2)

■

Exercise 8.1.12 Is $\otimes$ commutative? Is there an identity element? Does every element have an inverse?

8.1.2 Groups Defined

We have names to refer to algebraic structures with certain sets of features. Here is our first such name.

Definition 8.1.13 Suppose G is a set with associative binary operation $*$, identity element e, and with the property that every $g \in G$ has an inverse g^{-1} under $*$. Then the algebraic structure $(G, *, e, ^{-1})$ is called a *group*. If $*$ is a commutative binary operation, then G is called an *abelian* group (after the mathematician Neils Henrik Abel (1802–1829)). If G is finite, the cardinality of G is called the *order* of the group (usually denoted $o(G)$ instead of $|G|$).

According to Definition 8.1.13, a group has the following defining features.

(G1) The operation $*$ is well defined.

(G2) The operation $*$ is closed.

(G3) The operation $*$ is associative.

(G4) There is an identity element e.

(G5) Every element of G has an inverse under $*$.

Example 8.1.14 The real numbers with binary operation addition, identity zero, and additive inverses form an abelian group: $(\mathbb{R}, +, 0, -)$. For G1 is property A2, G2 is A3, G3 is A4, G4 is A6, and G5 is A7. That the group is abelian is property A5. Also, the integers and the rational numbers are abelian groups under addition. ■

Exercise 8.1.15 Determine whether each of the following is a group. If not, which of the properties G1–G5 fails to hold?

(a) The nonnegative real numbers with addition.

(b) The integers with multiplication.

(c) The nonzero real numbers with multiplication.

(d) For a non-empty set U, the power set of U with binary operation union.

(e) For a non-empty set U, the power set of U with binary operation intersection.

(f) For a non-empty set U, the power set of U with binary operation symmetric difference.[2]

Example 8.1.16 The algebraic structure $(S, \oplus, 0, -)$ from Example 8.1.9 is an abelian group. However, $(S, \otimes, 1, ^{-1})$ from Example 8.1.11 is not a group, for some elements do not have inverses. ■

[2] Perhaps you have finally earned the right to say, "Proof by picture" to prove associativity. Ask your instructor.

Example 8.1.17 Here is another example of a finite abelian group. In Exercise 2.2.5, you observed that $x^2 = -1$ has no real solution x. Nothing prevents us from creating a symbol, say i, declaring $i^2 = -1$, and then noting that i is not a real number. Now consider the set $S = \{\pm 1, \pm i\}$ with binary operation $\times$ defined according to Table 8.3. Then $(S, \times, 1, ^{-1})$ is an abelian group.

$\times$	1	-1	i	$-i$
1	1	-1	i	$-i$
-1	-1	1	$-i$	i
i	i	$-i$	-1	1
$-i$	$-i$	i	1	-1

(8.3)

■

Example 8.1.18 Similar to Example 8.1.17, define $\times$ on $Q = \{\pm 1, \pm i, \pm j, \pm k\}$ according to Table 8.4. The gist of this algebraic structure can be understood by noticing that $i^2 = j^2 = k^2 = -1$, but i, j, and k do not commute with each other. If you think of the letters i, j, and k being written on the face of a clock at 12, 4, and 8 o'clock, respectively, then multiplication of two different elements in the clockwise direction produces the third. Multiplication of two elements in the counterclockwise direction produces the negative of the third. For example, $j \times k = i$ and $i \times k = -j$. These three square roots of -1 are called *quaternions*, and they motivate a group with eight elements.

$\times$	1	-1	i	$-i$	j	$-j$	k	$-k$
1	1	-1	i	$-i$	j	$-j$	k	$-k$
-1	-1	1	$-i$	i	$-j$	j	$-k$	k
i	i	$-i$	-1	1	k	$-k$	$-j$	j
$-i$	$-i$	i	1	-1	$-k$	k	j	$-j$
j	j	$-j$	$-k$	k	-1	1	i	$-i$
$-j$	$-j$	j	k	$-k$	1	-1	$-i$	i
k	k	$-k$	j	$-j$	$-i$	i	-1	1
$-k$	$-k$	k	$-j$	j	i	$-i$	1	-1

(8.4)

■

EXERCISE 8.1.19 On the set $S = \{0, 1, 2, 3, 4, 5\}$, let $\circ$ be defined by Table 8.5.

(a) Explain how you know that $\circ$ is a binary operation.

(b) Is the operation commutative?

(c) What is the identity element?

(d) Determine the inverse of each element.

$\circ$	0	1	2	3	4	5
0	0	1	2	3	4	5
1	1	0	5	4	3	2
2	2	4	0	5	1	3
3	3	5	4	0	2	1
4	4	2	3	1	5	0
5	5	3	1	2	0	4

(8.5)

To verify that the binary operation in Exercise 8.1.19 is associative would be a formidable task if all you have is the Cayley table. Suffice it to say that $\circ$ is associative, so that S is a non-abelian group. This group is actually motivated by a group whose elements are functions. You will see this group in Section 8.4.

If someone gives you a set with a binary operation and asks you to show that it is a group, you must verify properties G1–G5. To give you an idea of how that might look, we will now walk through most of the details of a specific example.

Define the complex numbers by

$$\mathbb{C} = \{a + bi : a, b \in \mathbb{R}\} \tag{8.6}$$

At this point, i is a mere symbol with no meaning. We should think of $\mathbb{C}$ merely as a set of ordered real number pairs, the first of which stands alone and the second of which is tagged with an adjacent i. Furthermore, the $+$ in $a + bi$ is not meant to denote real addition, as if the expression $a + bi$ could be simplified. To avoid this possible confusion, some authors define elements of $\mathbb{C}$ as real number ordered pairs (a, b).

Next we define equality in $\mathbb{C}$, which for clarity we temporarily denote $\equiv$. We define $a_1 + b_1 i \equiv a_2 + b_2 i$, provided the real number equations $a_1 = a_2$ and $b_1 = b_2$ are satisfied. To show that $\equiv$ is an equivalence relation is pretty trivial. For example, to show $\equiv$ has property E2, suppose $a_1 + b_1 i \equiv a_2 + b_2 i$. Then $a_1 = a_2$ and $b_1 = b_2$. Since real number equality has property E2, we have $a_2 = a_1$ and $b_2 = b_1$. Thus $a_2 + b_2 i \equiv a_1 + b_1 i$.

Define the binary operation $\oplus$ in the following way.

$$(a + bi) \oplus (c + di) = (a + c) + (b + d)i \tag{8.7}$$

We claim that $\mathbb{C}$ is an abelian group under $\oplus$. Here are the steps of the proof in meticulous detail. Working through them will give you a good sense of direction in the exercise that follows, where you will show that the nonzero complex numbers $\mathbb{C}^{\times}$ with a form of multiplication is a group.

(G1) Suppose $a_1 + b_1 i \equiv a_2 + b_2 i$ and $c_1 + d_1 i \equiv c_2 + d_2 i$. Then $a_1 = a_2$, $b_1 = b_2$, $c_1 = c_2$, and $d_1 = d_2$. Since real number addition is well defined,

$a_1 + c_1 = a_2 + c_2$ and $b_1 + d_1 = b_2 + d_2$. Thus

$$\begin{aligned}(a_1 + b_1 i) \oplus (c_1 + d_1 i) &\equiv (a_1 + c_1) + (b_1 + d_1)i \\ &\equiv (a_2 + c_2) + (b_2 + d_2)i \\ &\equiv (a_2 + b_2 i) \oplus (c_2 + d_2 i)\end{aligned} \tag{8.8}$$

(G2) Pick $a + bi, c + di \in \mathbb{C}$. Since the real numbers are closed under addition, $a + c, b + d \in \mathbb{R}$, so that $(a + bi) \oplus (c + di) = (a + c) + (b + d)i \in \mathbb{C}$.

(G3)

$$\begin{aligned}[(a + bi) \oplus (c + di)] \oplus (e + fi) &= [(a + c) + (b + d)i] \oplus (e + fi) \\ &= [(a + c) + e] + [(b + d) + f]i \\ &= [a + (c + e)] + [b + (d + f)]i \\ &= (a + bi) \oplus [(c + e) + (d + f)i] \\ &= (a + bi) \oplus [(c + di) \oplus (e + fi)]\end{aligned} \tag{8.9}$$

(G4) Pick $a + bi \in \mathbb{C}$. Then $(a + bi) \oplus (0 + 0i) = (a + 0) + (b + 0)i = a + bi$. Similarly, $(0 + 0i) \oplus (a + bi) = (0 + a) + (0 + b)i = a + bi$. Thus $0 + 0i$ functions as an additive identity.

(G5) Pick $a + bi$. Then $a, b \in \mathbb{R}$, so that $-a, -b \in \mathbb{R}$ also. Thus $(-a) + (-b)i \in \mathbb{C}$. Furthermore, $(a + bi) \oplus [(-a) + (-b)i] = (a - a) + (b - b)i = 0 + 0i$, and $[(-a) + (-b)i] \oplus (a + bi) = (-a + a) + (-b + b)i = 0 + 0i$, so that $\mathbb{C}$ is closed under inverses.

Finally, $(\mathbb{C}, \oplus, 0 + 0i, -)$ is abelian because

$$\begin{aligned}(a + bi) \oplus (c + di) &= (a + c) + (b + d)i \\ &= (c + a) + (d + b)i \\ &= (c + di) \oplus (a + bi)\end{aligned} \tag{8.10}$$

If we had shown that $\oplus$ is commutative before we showed that properties G1–G5 hold, then the proofs of properties G4 and G5 would not have required the two-sided demonstrations we provided.

EXERCISE 8.1.20 Define a form of multiplication $\otimes$ on $\mathbb{C}^{\times}$ by

$$(a + bi) \otimes (c + di) \equiv (ac - bd) + (ad + bc)i \tag{8.11}$$

Show that $\mathbb{C}^{\times}$ is an abelian group under $\otimes$.[3,4]

EXERCISE 8.1.21 Prove that the integers with binary operation $a * b = a + b + 1$ form an abelian group.

EXERCISE 8.1.22 If we define a binary operation by $a * b = a + b - ab$, then there exists a real number x_0 such that $\mathbb{R} - \{x_0\}$ is a group under $*$. Find x_0 and show that $\mathbb{R} - \{x_0\}$ with $*$ is an abelian group.[5]

EXERCISE 8.1.23 Suppose $(G, *, e, ^{-1})$ is a group.

(a) Show that the identity element is unique.

(b) Prove the left and right cancellation laws:

(i) If $c * a = c * b$, then $a = b$.

(ii) If $a * c = b * c$, then $a = b$.

(c) Show that the inverse of a group element is unique.

(d) Show that $(a^{-1})^{-1} = a$ for all $a \in G$.

(e) Show that $(a * b)^{-1} = b^{-1} * a^{-1}$.

(f) Show inductively that $(a_1 * a_2 * \cdots * a_n)^{-1} = a_n^{-1} * a_{n-1}^{-1} * \cdots * a_1^{-1}$.

In Exercises 3.5.4 and 3.5.5, we defined a^n for nonzero real numbers a and all integers n, then derived rules for exponents. In the context of an arbitrary group G with binary operation $*$, we can now fix any $a \in G$ and make the same definitions for a^n, where we write

$$a^0 = e \tag{8.12}$$

$$a^{n+1} = a^n * a \quad \text{for } n \geq 0 \tag{8.13}$$

$$a^{-n} = (a^{-1})^n \tag{8.14}$$

Furthermore, by mimicking exactly the proofs from these exercises done in the context of the nonzero real numbers, we arrive at similar exponent rules for $*$ on G, except that one rule depends on G being abelian.

[3] Showing closure of $\otimes$ involves verifying that the product is never $0 + 0i$. Prove contrapositively by showing that if $(a + bi) \otimes (c + di) \equiv 0 + 0i$, then either $a = b = 0$ or $c = d = 0$, so that either $a + bi$ or $c + di$ is not an element of $\mathbb{C}^{\times}$.

[4] If $(a + bi) \otimes (c + di) \equiv 0 + 0i$, then $ac - bd = 0$ and $ad + bc = 0$. Square these and add.

[5] You can't find x_0 until you have found the identity.

Theorem 8.1.24 Let G be a group, $a, b \in G$, and let m and n be integers. Then

$$a^m * a^n = a^{m+n} \tag{8.15}$$

$$(a^m)^n = a^{mn} \tag{8.16}$$

Furthermore, if G is abelian,

$$(a * b)^n = a^n * b^n \tag{8.17}$$

Exercise 8.1.25 Suppose $(G, *, e, ^{-1})$ is a group such that $a * a = e$ for all $a \in G$. Then G is abelian.[6]

Exercise 8.1.26 Let $(G, *, e, ^{-1})$ be a group, and fix some $g \in G$. Define $f : G \to G$ by $f(x) = g * x$ for all $x \in G$. Then f is a bijection from G to itself.

Convenience is desirable if it costs nothing in clarity. For this reason, we often refer to a group $(G, *, e, ^{-1})$ simply as group G, and we sometimes use juxtaposition of terms ab to indicate the binary operation $a * b$, just as we do with multiplication in the real numbers. If we are working with two groups G_1 and G_2, we should probably be more careful at first to distinguish between the symbols for the binary operations, and we might denote the identity elements as e_1 and e_2, respectively, just to be clear.

8.2 Subgroups

You might have noticed that S from Example 8.1.17 is a subset of Q from Example 8.1.18. More than that, the binary operation on S is the same as that on Q, in the sense that Table 8.3 is a subtable of Table 8.4. Thus S is a subset of Q that is closed under the binary operation, contains the identity, and is closed under inverses. We give such a subset a name.

8.2.1 Subgroups Defined

Definition 8.2.1 Suppose $(G, *, e, ^{-1})$ is a group, H is a subset of G, and H is a group under the same operation. Then $(H, e, *, ^{-1})$ is called a *subgroup* of G, and we write $H < G$. If H is a proper subset of G, then H is called a *proper subgroup* of G.

If we are given a subset H of a group $(G, *, e, ^{-1})$ and we are asked to show that H is a subgroup of G, we must show that H is itself a group under $*$. Properties

[6] $a * a = e$ is equivalent to $a = a^{-1}$. Use Exercise 8.1.23(e).

G1 and G3 are automatically satisfied on H, for if $*$ is well defined and associative on all of G, then certainly it is well defined and associative when restricted to H. We say that H *inherits* these properties from G. We must therefore demonstrate only the remaining properties.

(H1) The operation is closed on H (G2).

(H2) The identity element is in H (G4).

(H3) The inverse of every element of H is also in H; that is, H is closed under inverses (G5).

Example 8.2.2 For any group G, $\{e\}$ and G itself satisfy properties H1–H3 and are therefore subgroups of G. We call $\{e\}$ the *trivial subgroup*. ■

Example 8.2.3 The set of even integers is a subgroup of $(\mathbb{Z}, +, 0, -)$ since it is closed under addition, contains zero, and is closed under negation. ■

Example 8.2.3 is a special case of the next result.

EXERCISE 8.2.4 Let n be an integer and consider $S = \{kn : k \in \mathbb{Z}\}$, the set of all integer multiples of n. Show that S is a subgroup of the integers under addition.

EXERCISE 8.2.5 Find all subgroups of each of the following groups.

(a) The group in Example 8.1.9

(b) The quaternion group in Example 8.1.18

(c) The group with Cayley Table 8.18

$\oplus$	0	1	2	3	4
0	0	1	2	3	4
1	1	2	3	4	0
2	2	3	4	0	1
3	3	4	0	1	2
4	4	0	1	2	3

(8.18)

EXERCISE 8.2.6 From the multiplicative group in Exercise 8.1.20, let

$$H = \{a + bi \in \mathbb{C} : a^2 + b^2 = 1\} \tag{8.19}$$

Show that H is a subgroup of $\mathbb{C}^\times$.

EXERCISE 8.2.7 Let $G = \{a + b\sqrt{2} : a, b \in \mathbb{Q},\ a \text{ and } b \text{ not both zero}\}$. Clearly, G is a subset of the real numbers. Also, by Exercise 3.10.6, $a + b\sqrt{2} = c + d\sqrt{2}$ if and only if $a = c$ and $b = d$. We want to show that G is a subgroup of $\mathbb{R}^\times$, the multiplicative group of nonzero real numbers.

(a) Explain why zero is not in G.

(b) Show that G is a subgroup of $(\mathbb{R}^{\times}, \times, 1, ^{-1})$ by showing it has properties H1–H3.

EXERCISE 8.2.8 Define the *center* of a group G to be the set of all elements that commute with all elements of G. That is, the center is

$$\{a \in G : a * x = x * a \text{ for all } x \in G\} \tag{8.20}$$

Show that the center of G is a subgroup of G.

EXERCISE 8.2.9 Suppose $\{H_\alpha\}_{\alpha \in \mathcal{A}}$ is a family of subgroups of a group G. Determine with proof whether each of the following is a subgroup of G.

(a) $\bigcap_{\alpha \in \mathcal{A}} H_\alpha$

(b) $\bigcup_{\alpha \in \mathcal{A}} H_\alpha$

Suppose H is a subgroup of G and $H \subseteq H_1 \subseteq G \subseteq G_1$. We can make the following observations about relationships between these sets. First, if H_1 is a group under the same binary operation, then the fact that H is a subgroup of G implies H is a subgroup of H_1 also. For if H exhibits properties H1–H3 as a subset of G, it also does so as a subset of H_1. Similarly if G_1 is a group, then H is a subgroup of G_1. The most efficient way to describe this latter relationship is to say that $<$ is transitive: If $H < G$ and $G < G_1$, then $H < G_1$.

8.2.2 Generated Subgroups

From the quaternion group Q in Example 8.1.18, let's take some subset of Q, say $A = \{i, j\}$. Although A is a subset of Q, A is clearly not a subgroup of Q. We can, however, determine the smallest subgroup of Q that contains all elements of A. We now address the existence and possible uniqueness of what we call the *subgroup generated by* A. Let's begin by defining the term that conceptually means the smallest subgroup containing all elements of a given subset of the group.

Definition 8.2.10 Suppose G is a group and A is a non-empty subset of G. Suppose also that H is a subset of G with the following properties.

(U1) $A \subseteq H$.

(U2) H is a subgroup of G.

(U3) If B is a subgroup of G, and $A \subseteq B$, then $H \subseteq B$.

Then H is called a *subgroup generated by* A, and is denoted (A).

Notice how properties U1–U3 lay out the appropriate criteria by which a set qualifies as a smallest subgroup that contains all elements of A. Property U1 guarantees that any set that we would call (A) does in fact contain all elements of A, and property U2 guarantees that it is indeed a subgroup of G. Property U3 guarantees that no other subgroup of G that contains all elements of A will be any smaller. The fact that (A) exists uniquely is guaranteed by the next exercise, which gives us one way to visualize its construction.

Let A be a subset of a group G, and consider the family of subgroups of G, where each set in the family is a superset of A. That this family is non-empty is immediate, for G itself is a subgroup of G and is a superset of A. In the next exercise, you will show that the subgroup generated by A exists uniquely. What you will show is that the intersection of all subgroups of G that are supersets of A has properties U1–U3, and that any two subsets of G that both have properties U1–U3 are actually the same.

EXERCISE 8.2.11 Let G be a group, and suppose A is a non-empty subset of G. Then (A) exists uniquely. In fact,

$$(A) = \bigcap_{\substack{J<G \\ J\supseteq A}} J \tag{8.21}$$

Equation (8.21) might not be the best way to construct (A). It might be easier to begin with the elements of A and build up (A) by tossing in only the necessary elements of G until you're sure you're finished.

Example 8.2.12 From Example 8.1.18, and $A = \{i, j\}$, determine (A).

Solution Clearly, (A) must contain 1, and by property H1, it must contain $i^2 = -1$ and $i * j = k$. Appealing to H3, we must have $-i, -j, -k \in (A)$ also. Thus $(A) = Q$. ■

EXERCISE 8.2.13 For the group from Example 8.1.9, determine the subgroup generated by the following sets.

(a) $\{2, 4\}$

(b) $\{2, 3\}$

(c) $\{5\}$

8.2.3 Cyclic Subgroups

The subgroup generated by a singleton set $\{a\}$ is usually denoted (a) instead of $(\{a\})$. In this case, we call (a) the subgroup of G generated by the element a, and a is called a *generator* of this subgroup. Any subgroup that has a generator is called a *cyclic subgroup*. The easiest way to see what (a) looks like is to build it from the bottom up, so to speak, and then demonstrate that what you have created satisfies U1–U3. Using the definition of a^n in Eqs. 8.12–8.14, consider the set

$$S = \{a^n : n \in \mathbb{Z}\} \tag{8.22}$$

The claim is that $S = (a)$, which we verify here by showing that S has properties U1–U3, thereby earning the right to be called (a). The details will help you in Exercise 8.2.24.

Let $n = 1$ to see that $a \in S$, so that S has property U1. To show S has property U2, we must show it has properties H1–H3.

(H1) Pick $x, y \in S$. Then there exist integers m and n such that $x = a^m$ and $y = a^n$. Now $m + n$ is also an integer, so $a^m * a^n = a^{m+n} \in S$. Thus S is closed under $*$.

(H2) Letting $n = 0$, we have that $e = a^0 \in S$.

(H3) Pick $a^n \in S$. Since Theorem 8.1.24 applies for all integers n, we have that $(a^n)^{-1} = a^{-n} \in S$. Thus S is closed under inverses.

To show that S has property U3, suppose B is a subgroup and $a \in B$. We show $S \subseteq B$ by showing $a^n \in B$ for all n. Since $a \in S$ and B is closed under $*$, it must be that $a^n \in B$ for all positive integers n. Certainly, $a^0 = e \in B$, and since B is closed under inverses, $a^{-n} \in B$ for all positive integers n. Thus $S \subseteq B$, and we have finished the proof that $(a) = \{a^n : n \in \mathbb{Z}\}$.

Example 8.2.14 In the multiplicative group of nonzero real numbers,

$$(3) = \{3^n : n \in \mathbb{Z}\} = \{\ldots, 1/27, 1/9, 1/3, 1, 3, 9, 1/27, \ldots\} \quad \blacksquare$$

Example 8.2.15 We construct (3) in the *additive* group $(\mathbb{Z}, +, 0, -)$. Note that zero is the identity in this context, so that $3^0 = 0$. Constructing 3^n for $n \geq 2$ is repeated addition, not repeated multiplication. Thus

$$3^1 = 3, \quad 3^2 = 3 + 3 = 6, \quad 3^3 = 6 + 3 = 9, \quad 3^4 = 9 + 3 = 12, \quad \text{etc.} \tag{8.23}$$

Constructing 3^{-n} involves negation, not reciprocation, so that

$$\begin{aligned} 3^{-1} &= -3 \\ 3^{-2} &= (3^{-1})^2 = [(-3) + (-3)] = -6 \\ 3^{-3} &= (3^{-1})^3 = [(-6) + (-3)] = -9, \quad \text{etc.} \end{aligned} \tag{8.24}$$

Thus $(3) = \{3k : k \in \mathbb{Z}\}$. $\blacksquare$

Example 8.2.15 suggests it might be worthwhile to rewrite the exponent definitions in Eqs. (8.12)–(8.14) and the exponent rules in Eqs. (8.15)–(8.17) in what we call their *additive form*. In particular, if $(G, +, 0, -)$ is a group and $a \in G$, the fact that the binary operation is a form of addition suggests that, instead of writing $a^0 = e$, which seems to connote multiplication, we instead write

$$0a = 0 \tag{8.25}$$

We must be careful, though, for the zero on the left-hand side of Eq. (8.25) is an integer, whereas the zero on the right is the identity element of the group. Saying $0a = 0$ means that the group element a is, in a sense, added to itself zero times, not multiplied by itself, and this is to be defined as the zero (identity) element of the group. Similarly, the additive form of Eq. (8.15) would be

$$ma + na = (m + n)a \tag{8.26}$$

EXERCISE 8.2.16 Let $(G, +, 0, -)$ be an additive group. Write Eqs. (8.12)–(8.17) in their additive form.

EXERCISE 8.2.17 Determine the following cyclic subgroups.

(a) (5) in $(\mathbb{R}^\times, \times, 1, ^{-1})$

(b) (5) in $(\mathbb{Z}, +, 0, -)$

(c) (5) in the group from Example 8.1.9

(d) (2) in the group from Example 8.1.9

(e) (3) in the group from Example 8.1.9

(f) (0) in the group from Example 8.1.9

(g) (a), where a is any element other than the identity in the group with Cayley Table 8.18

(h) (j) in the quaternion group from Example 8.1.18

(i) $(2i)$ in $(\mathbb{C}^\times, \times, 1, ^{-1})$

(j) $(\frac{1}{\sqrt{2}} + \frac{1}{\sqrt{2}}i)$ in $(\mathbb{C}^\times, \times, 1, ^{-1})$

As Exercise 8.2.17 illustrates, even though a group is infinite, the subgroup generated by an element might or might not be infinite. In Example 8.2.14, the subgroup generated by 3 is an infinite cyclic subgroup. But let's consider the subgroup generated by i in $(\mathbb{C}^\times, \{0 + 0i\}, \times, 1 + 0i, ^{-1})$. First, $i^0 = 1$, $i^1 = i$, $i^2 = -1$ and $i^3 = -i$. But $i^4 = 1$, and 4 is the smallest positive power of i for which this is true. Furthermore, for any integer n, the division algorithm allows us to write $n = 4k + r$ where $0 \le r \le 3$, so that $i^n = i^{4k+r} = (i^4)^k \cdot i^r = i^r$. Thus every power of i can be simplified to an element of $\{i^0, i^1, i^2, i^3\}$, and we have shown that

$$(i) = \{i^n : n \in \mathbb{Z}\} = \{i^0, i^1, i^2, i^3\} = \{1, i, -1, -i\} \tag{8.27}$$

so that (i) is a finite cyclic subgroup in an infinite group.

This suggests a general result for some cyclic subgroups. Suppose G is a group and $a \in G$. Suppose there exists some integer k such that $a^k = e$. Since a^{-k} would also be the identity, we may assume $k \geq 1$. Of all positive values of k for which $a^k = e$, let n be the smallest and consider the set

$$T = \{a^0, a^1, a^2, \ldots, a^{n-1}\} = \{a^k : 0 \leq k \leq n-1\} \tag{8.28}$$

Since n is the smallest positive integer for which $a^n = e$, no a^k is the identity for any $1 \leq k \leq n-1$. Thus e appears exactly one time in Eq. (8.28) as a^0. Furthermore, all elements of T are distinct.

EXERCISE 8.2.18 Show that the elements of T in Eq. (8.28) are distinct by showing that if $0 \leq k < l \leq n-1$, then $a^k \neq a^l$ (or the contrapositive).

Clearly, $\{a^k : 0 \leq k \leq n-1\} \subseteq \{a^k : k \in \mathbb{Z}\}$, so that $T \subseteq (a)$.

EXERCISE 8.2.19 Show $T \supseteq (a)$ to have $T = (a)$.[7]

Thus we have shown the following.

Theorem 8.2.20 Let G be a group and let $a \in G$. Suppose $a^n = e$ for some positive integer n, and suppose n is the smallest such positive integer. Then

$$(a) = \{e, a, a^2, a^3, \ldots, a^{n-1}\} \tag{8.29}$$

If n is the smallest positive integer for which $a^n = e$, we say that the element a has *order* n, and we denote the order of a by $o(a)$. If $a^n \neq e$ for all positive integers n, we say a has *infinite order*. In Definition 8.1.13, we defined the order of a group as its cardinality. Here we are defining the order of an element of a group in terms of its powers. By constructing (a) as we have done here, we have demonstrated that these two uses of the term order are tied together.

Theorem 8.2.21 If G is a group and $a \in G$ has order n, then the subgroup generated by a has order n. Conversely, if (a) has n elements, then a has order n.

With Theorem 8.2.21, we can use $o(a)$ to mean either the order of the element a or the order of the subgroup generated by a. Naturally, if G is finite and $a \in G$, the subgroup generated by a will be finite.

Example 8.2.22 In the quaternion group from Example 8.1.18,

$$(k) = \{k^n : 0 \leq n \leq 3\} = \{1, k, -1, -k\}$$

so that $o(k) = 4$. ■

[7] Pick $a^k \in (a)$ and apply the division algorithm to k and n.

EXERCISE 8.2.23 If G is a group and $a \in G$, then (a) is abelian.

EXERCISE 8.2.24 Suppose G is an abelian group, and $a, b \in G$. Then

$$(\{a, b\}) = \{a^m b^n : m, n \in \mathbb{Z}\} \tag{8.30}$$

EXERCISE 8.2.25 Write Eq. (8.30) in its additive form.

EXERCISE 8.2.26 Consider the group $(\mathbb{Z}, +, 0, -)$, and let a and b be nonzero integers and $g = \gcd(a, b)$. Then $(\{a, b\}) = (g)$.

EXERCISE 8.2.27 Suppose G is a group and $a \in G$. Let m and n be positive integers, and suppose $\gcd(m, n) = g$. Then $(\{a^m, a^n\}) = (a^g)$.

EXERCISE 8.2.28 For the additive group of integers, determine the following without proof.

(a) $(4) \cap (6)$

(b) $(10) \cap (3)$

(c) $(8) \cap (16)$

(d) $(12) \cap (16) \cap (28)$

(e) $(\{4, 6\})$

(f) $(\{4, 7\})$

(g) $(\{12, 16, 28\})$

If G is a group and there exists some element g such that $(g) = G$, then G is said to be a *cyclic* group and g is called a *generator* of the group. There are some nifty little theorems about cyclic groups. For example, by Exercise 8.2.23, a cyclic group is abelian. Here is another.

EXERCISE 8.2.29 A cyclic group is countable.

EXERCISE 8.2.30 Determine with explanation whether each of the following groups is cyclic.

(a) $(\mathbb{Z}, +, 0, -)$

(b) The group from Example 8.1.9

(c) $(\mathbb{R}^\times, \times, 1, ^{-1})$

(d) The quaternion group from Example 8.1.18

(e) $(\mathbb{Q}, +, 0, -)$

8.3 Quotient Groups

In this section, we return to the group in Example 8.1.9 to derive it in a formal way from the integers. The result of this construction is called the *integers modulo n* ($n = 6$ in Example 8.1.9). This construction serves as a good illustration of a *quotient group*, which we will discuss as a generalization of the process of deriving the integers modulo n.

8.3.1 Integers Modulo n

In Example 8.1.9, we noted that $(S, \oplus, 0, -)$ is a group, where $S = \{0, 1, 2, 3, 4, 5\}$ and $\oplus$ was defined in Table 8.1. You might have noticed the similarity of $\oplus$ to regular addition of integers, except that it was a sort of circular addition. Any sum such as $4 + 4$ that exceeds 5 is reduced by 6 to guarantee that the sum is in S. This circular summing is sometimes called clock arithmetic, where in this case, the numbers $\{0, 1, 2, 3, 4, 5\}$ can be written around the perimeter of a clock and addition can be performed in a natural circular way. This example of a group might sound a bit simplistic but actually, in spite of its simplicity, is a most important example in group theory. Perhaps we should say *because of* its simplicity, its importance in group theory is particularly striking and beautiful.

Here is a standard way of constructing this group by beginning with the integers. It takes a few steps and might seem a bit esoteric, but it is representative of a more general procedure. Study it carefully, because some of the details are left to you, and you will mimic the details as you prove the corresponding results of the more general procedure in an exercise to follow. Instead of considering the special case of $S = \{0, 1, 2, 3, 4, 5\}$, we consider a more general case $\{0, 1, 2, \ldots, n - 1\}$.

Step 1: Construct the set. Start with the group $(\mathbb{Z}, +, 0, -)$, and let n be a given positive integer. Let

$$(n) = \{kn : k \in \mathbb{Z}\} = \{\ldots, -3n, -2n, -n, 0, n, 2n, 3n, \ldots\} \tag{8.31}$$

be the subgroup of the integers generated by n. Define an equivalence relation on the integers as follows. Define

$$x \equiv_n y \quad \text{if and only if} \quad x - y \in (n) \tag{8.32}$$

That is, $x \equiv_n y$ provided $x - y = kn$ for some integer k. This is the same definition of equivalence that we used in Exercise 3.6.18, so proving $\equiv_n$ is an equivalence relation on the integers has already been done. Recall from Exercise 3.6.23 that $x \equiv_n y$ if and only if x and y have the same remainder when divided by n according to the division algorithm. Also, recall the equivalence classes that arise from this definition of equivalence:

$$\begin{aligned}
[0] &= \{\ldots, -2n, -n, 0, n, 2n, \ldots\} \\
[1] &= \{\ldots, -2n+1, -n+1, 1, n+1, 2n+1, \ldots\} \\
[2] &= \{\ldots, -2n+2, -n+2, 2, n+2, 2n+2, \ldots\} \\
&\vdots \\
[k] &= \{\ldots, -2n+k, -n+k, k, n+k, 2n+k, \ldots\} \\
&\vdots \\
[n-1] &= \{\ldots, -n-1, -1, n-1, 2n-1, 3n-1, \ldots\}
\end{aligned} \tag{8.33}$$

Notice the very important fact that $[0] = (n)$; that is, the equivalence class of the identity in $(\mathbb{Z}, +, 0, -)$ is the subgroup from which the equivalence is defined. Lump these n equivalence classes into a family and call it $\mathbb{Z}/(n)$, the *integers modulo n*.

$$\mathbb{Z}/(n) = \{[0], [1], [2], \ldots, [n-1]\} \tag{8.34}$$

Recall that every equivalence class has infinitely many names, depending on the representative element by which we choose to address it. For example, in $\mathbb{Z}/(6)$, $[5] = [-1] = [41]$.

Step 2: Define a binary operation on the created set. On the elements of $\mathbb{Z}/(n)$, which are themselves sets, we define a form of addition $\oplus_n$ in the following way:

$$[a] \oplus_n [b] = [a +_{\mathbb{Z}} b] \tag{8.35}$$

where $+_{\mathbb{Z}}$ denotes the sum of the representative integers a and b. For example, $[4] \oplus_6 [5] = [4+5] = [9] = [3]$, where we choose to refer to the sum as [3] since [3] is the same set as [9] and 3 is a representative element between 0 and 5.

Notice that this new binary operation $\oplus_n$ is a way of combining two sets to produce another set. It is not union or intersection, which up to now are the only binary operations on sets we have seen. Instead, $\oplus_n$ combines $[a]$ and $[b]$ by using representative elements from each to produce a representative element of the set we are defining to be $[a] \oplus_n [b]$.

Step 3: Show that the created set with its binary operation gives rise to a group. We must show that $\oplus_n$ is well defined, closed, and associative. We must also find an identity element and find inverses for all elements. Clearly, $\oplus_n$ is closed. For if $[a], [b] \in \mathbb{Z}/(n)$, then $a + b$ is an integer. Since the equivalence classes in $\mathbb{Z}/(n)$ partition the integers, $a + b$ is in some equivalence class. Thus $[a + b] \in \mathbb{Z}/(n)$.

Exercise 8.3.1 Finish the proof that $\mathbb{Z}/(n)$ with operation $\oplus_n$ is a group with the following steps.

(a) Show $\oplus_n$ is well defined.

(b) Show $\oplus_n$ is associative.

(c) Determine the identity element.

(d) For an arbitrarily chosen $[k]$, determine the inverse of $[k]$.

And we are done. Table 8.36 displays the final product for the case $\mathbb{Z}/(6)$.

$\oplus_n$	[0]	[1]	[2]	[3]	[4]	[5]
[0]	[0]	[1]	[2]	[3]	[4]	[5]
[1]	[1]	[2]	[3]	[4]	[5]	[0]
[2]	[2]	[3]	[4]	[5]	[0]	[1]
[3]	[3]	[4]	[5]	[0]	[1]	[2]
[4]	[4]	[5]	[0]	[1]	[2]	[3]
[5]	[5]	[0]	[1]	[2]	[3]	[4]

(8.36)

Let's take a closer look at the equivalence classes $[0], [1], \ldots, [n-1]$, for there is another standard notation for these sets that we will use below when we derive quotient groups in general. Though we know from Eq. (8.33) what is in each equivalence class, let's consider a way to visualize an arbitrary equivalence class $[k]$ in terms of the subgroup of the integers $(n) = [0]$ that motivated this whole mess in the first place.

Picture all the integers on the number line in standard fashion, but pretend that each integer is like a key on a piano keyboard. Take countably infinitely many of your friends, and each of you place a finger on the elements of $(n) = [0]$, so that you are pointing to all the multiples of n: $\{\ldots, -2n, -n, 0, n, 2n, \ldots\}$. Now suppose you want to point to all the elements of some equivalence class $[k]$. How can you do it? Everyone in unison should lift his or her finger off the keyboard, and everyone should shuffle over to the right k units, then put his or her finger back down. In other words, to generate all the elements of $[k]$, take all the elements of (n) and add k to each one, so that you are translating the entire set (n) through the integers by k units. Notice in this imagery, there will always be some person pointing to one of the integers $\{0, 1, 2, \ldots, n-1\}$. Here is another way to describe what we have done.

$$[k] = \{x + k : x \in (n)\} \tag{8.37}$$

Let's create new notation for the construction in Eq. (8.37), writing

$$(n) + k = \{x + k : x \in (n)\} \tag{8.38}$$

This is a slight abuse of orthodox notation because it appears that we are combining the subgroup (n) with an integer k using integer addition. This notation is standard, however, and it makes the definition of addition in Eq. (8.35) look like this:

$$\big((n)+a\big)\oplus_n\big((n)+b\big)=(n)+\big(a+b\big) \tag{8.39}$$

Take note of the multiple uses of parentheses in Eq. (8.39). Some are used to denote a generated subgroup, and others are used merely to group symbols. With this notation for the equivalence classes in $\mathbb{Z}/(n)$, the Cayley table, written in its painfully rigorous form, would look like Table 8.40 for the case $n=6$. Once again, notice that every coset can be addressed by a unique integer $0,\ldots,5$.

$\oplus_n$	$(n)+0$	$(n)+1$	$(n)+2$	$(n)+3$	$(n)+4$	$(n)+5$
$(n)+0$	$(n)+0$	$(n)+1$	$(n)+2$	$(n)+3$	$(n)+4$	$(n)+5$
$(n)+1$	$(n)+1$	$(n)+2$	$(n)+3$	$(n)+4$	$(n)+5$	$(n)+0$
$(n)+2$	$(n)+2$	$(n)+3$	$(n)+4$	$(n)+5$	$(n)+0$	$(n)+1$
$(n)+3$	$(n)+3$	$(n)+4$	$(n)+5$	$(n)+0$	$(n)+1$	$(n)+2$
$(n)+4$	$(n)+4$	$(n)+5$	$(n)+0$	$(n)+1$	$(n)+2$	$(n)+3$
$(n)+5$	$(n)+5$	$(n)+0$	$(n)+1$	$(n)+2$	$(n)+3$	$(n)+4$

(8.40)

Now let's relax the notation for the specific context of $\mathbb{Z}/(n)$. First, $\mathbb{Z}/(n)$ is usually written $\mathbb{Z}_n$. Also, instead of always writing $(n)+k$, we allow ourselves simply to write k, understanding that we are not talking about a single integer but the entire equivalence class of integers of which k is a representative element. We also generally revert to the regular addition sign $+$ rather than writing $\oplus_n$. With this simplified notation, Table 8.40 becomes Table 8.1.

This all boils down to the fact that mathematicians get a little lazy and let the familiar notation from $(\mathbb{Z},+,0,-)$ also serve for the group $\mathbb{Z}_n$ with addition, so that $(\mathbb{Z}_n,+,0,-)$ is a relaxed notation for $(\mathbb{Z}/(n),\oplus_n,(n)+0,-)$. This is probably all right because we are less interested in fancy notation than we are in the internal structure of $\mathbb{Z}_n$ as a group. And it is just as helpful to visualize $\mathbb{Z}_n$ as clock arithmetic with regular old numbers as it is to think in terms of combining equivalence classes of integers. Understood this way, we can then write something like $17+9=_6 2$ and know just what we mean.

8.3.2 Quotient Groups

The derivation of the integers modulo n is affected by the fact that integer addition is commutative. Did you notice at what point in Exercise 8.3.1 you exploited this fact?[8] Now let's generalize the program for deriving $\mathbb{Z}_n$ to create a quotient group in a more abstract setting. Before we get started, however, we must explain a somewhat restricted approach we will take. At first, we are going to restrict ourselves to an abelian group as a starting point, which we should not have to do.

[8] Somewhere in showing $\oplus_n$ is well defined.

We will allude to the reason for this at the end of the discussion, but for now, do not use the fact that the group is abelian anywhere it is not absolutely necessary, and take note of the step(s) in the proof where you do use it. The proofs of all the steps along the way in this derivation are left to you as exercises. As you read them, refer to their parallel steps in $\mathbb{Z}_n$, and note how the new notation here is a generalization of the $\mathbb{Z}/(n)$ notation to an arbitrary group.

Step 1: Given an *abelian* group $(G, *, e, ^{-1})$ and a subgroup H, we define a relation $\equiv_H$ on G in the following way. Define

$$a \equiv_H b \Leftrightarrow a * b^{-1} \in H \tag{8.41}$$

By the next exercise, $\equiv_H$ is an equivalence relation, so that G is partitioned into equivalence classes. We denote the set of all equivalence classes G/H, which is read "G mod H."

EXERCISE 8.3.2 Show that $\equiv_H$ is an equivalence relation.

Before we continue the construction of the group G/H by defining its binary operation, let's see what the equivalence classes generated by the equivalence definition in (8.41) look like. In $\mathbb{Z}_n$ we have a pretty clear picture of what is in each equivalence class, as illustrated by Eq. (8.33). In a general group, however, we want a way to visualize the elements of $[g]$ for a given $g \in G$. In the next exercise, you will show that $x \in [g]$ if and only if there exists some $h \in H$ such that $x = h * g$.

EXERCISE 8.3.3 Show that $[g] = \{h * g : h \in H\}$.

Exercise 8.3.3 says that we can think of the elements of $[g]$ as the result of taking all elements of H and using the binary operation[9] to scoot them all over g "amount." The set $[g]$ is called the *right coset of H generated by g*, and instead of being written in equivalence class notation is denoted

$$H * g = \{h * g : h \in H\} \tag{8.42}$$

A coset is strikingly similar to the construction in (8.37) and can be visualized in a similar way as a translation of all elements of H through G by g units, so to speak. We then write $G/H = \{H * g : g \in G\}$, the family of all cosets of H that can be created by letting g take on all values in G.

We are now ready to continue the construction of the quotient group by defining a binary operation on G/H.

[9] With g on the right of each $h \in H$, which will matter later.

Step 2: Define an operation $*_H$ on G/H in the following way.

$$(H * a) *_H (H * b) = H * (a * b) \tag{8.43}$$

Note the distinction between the uses of $*$ in this definition. We are using $*_H$ to denote the operation we are defining on the family of cosets of H. The $*$ in $H * a$ is the standard notation for the coset generated by a, comparable to $(n) + k$ in Eq. (8.38). Finally, the $*$ in $(a * b)$ is the binary operation on G.

Step 3: Show that $*_H$ is a well-defined, closed, associative binary operation on G/H, find the identity element, and find inverses for elements.

Exercise 8.3.4 Show that G/H with operation $*_H$ is a group with the following steps.

G1: Show $*_H$ is well defined on G/H.

G2: Show $*_H$ is closed on G/H.

G3: Show $*_H$ is associative on G/H.

G4: Show $H * e$ is the identity element of G/H.

G5: Show that every $H * a$ has an inverse in G/H.

Unless you are unnecessarily sloppy with your algebraic manipulation in completing steps 1–3, there is only one place you must exploit the fact that G is abelian, and it is the same place where you exploited commutativity of integer addition in your work with the integers modulo n. In Section 8.5, we will return to this construction in the context of an arbitrary group that might not be abelian. But even then, if the program is going to work, we cannot completely do without some way to switch the order of certain elements. In Section 8.5, we will guarantee the property we need by requiring H to be a special kind of subgroup called a *normal subgroup*. The relationship of a normal subgroup to the entire group will look a bit like commutativity, but weaker, so that the construction of the quotient group can be realized in as general a group as possible.

The set of all cosets with its binary operation forms the new group, which we call the *quotient group* generated by H. Here is the definition all in one piece.

Definition 8.3.5 Suppose $(G, *, e, ^{-1})$ is a group (abelian), and let H be a subgroup of G. Define an equivalence relation $\equiv_H$ by declaring $a \equiv_H b$ provided $a * b^{-1} \in H$. Then

$$G/H = \{H * g : g \in G\} \tag{8.44}$$

with binary operation $*_H$ defined by

$$(H * a) *_H (H * b) = H * (a * b) \tag{8.45}$$

identity element $H * e = H$, and inverses $(H * a)^{-1} = H * a^{-1}$ is called the *quotient group* of G created by *modding out* the subgroup H. Elements of G/H are called *right cosets* of H.

For equivalence classes in general, we know that $[a] = [b]$ if and only if a and b are equivalent; otherwise $[a]$ and $[b]$ are disjoint. In the context of Definition 8.3.5, these facts become $H * a = H * b$ if and only if $a * b^{-1} \in H$; otherwise $H * a$ and $H * b$ are disjoint. In other words, a and b generate the same coset of H if and only if $a * b^{-1} \in H$. Otherwise the cosets they generate are disjoint.

Example 8.3.6 Let G be the set of all functions $f : \mathbb{R} \to \mathbb{R}$. For two functions f and g, we define their sum $f + g$ by the rule $[f + g](x) = f(x) + g(x)$. Clearly, the function $\mathbf{0}$ defined by $\mathbf{0}(x) = 0$ for all x is the identity element, and $-f$ is the function defined by $[-f](x) = -f(x)$. Thus $(G, +, \mathbf{0}, -)$ is an abelian group. Let H be the subset of G consisting of all constant functions, which is clearly a subgroup of G.

What do the elements of the quotient group G/H look like? If f is a given function in G, then $H + f = \{h + f : h \in H\}$ is the set of all translations of f up and down in the xy-plane by the constant functions in H. That is, $g \in H + f$ if and only if there exists some constant function h such that $g = h + f$. Equivalently, $g \in H + f$, provided $g - f \in H$, which means that $g - f$ is a constant function. This might remind you of a fact from calculus. If f is any antiderivative of a function f_1, then $H + f$ is the set of all the antiderivatives of f_1. If you like, every coset can be addressed by choosing a representative element, perhaps the function in the coset that passes through the origin. ■

Exercise 8.3.7 Since the real numbers under addition form an abelian group with the integers as a subgroup, we may discuss $\mathbb{R}/\mathbb{Z}$.

(a) In constructing $\mathbb{R}/\mathbb{Z}$, what is the definition of equivalence that gives rise to the right cosets of $\mathbb{Z}$?

(b) Construct the following cosets by listing some of their elements.

(i) $\mathbb{Z} + 1.4$
(ii) $\mathbb{Z} + \sqrt{2}$
(iii) $\mathbb{Z} + 5$

(c) What is a convenient subset of the real numbers from which we may choose a unique representative element of each coset?[10]

(d) Evaluate the following in $\mathbb{R}/\mathbb{Z}$. Express your answer as a coset $\mathbb{Z} + r$, where r is an element of your answer to part (c).

(i) $(\mathbb{Z} + 2.2) +_{\mathbb{Z}} (\mathbb{Z} + 8.14)$
(ii) $(\mathbb{Z} + 3.8) +_{\mathbb{Z}} (\mathbb{Z} + 0)$

[10] There are infinitely many answers, but one of them is considered standard.

(iii) $(\mathbb{Z}+3.8)+_{\mathbb{Z}}(\mathbb{Z}+1.2)$

(iv) $-(\mathbb{Z}+11.23)$

(v) $(\mathbb{Z}+3.8)-_{\mathbb{Z}}(\mathbb{Z}+1.2)$

(e) Describe addition in $\mathbb{R}/\mathbb{Z}$ with an analogy similar to the clock arithmetic used to describe addition in the integers modulo n.

8.3.3 Cosets and Lagrange's Theorem

Even if G is not an abelian group and the family of cosets of a subgroup of G do not give rise to a quotient group, we can still derive some important results about elements and subgroups of G by looking at the cosets of the subgroup. One thing to keep in mind if G is not abelian is that we have to distinguish between left and right cosets. For a subgroup H and some fixed $g \in G$, we use the notation

$$g * H = gH = \{g * h : h \in H\} \tag{8.46}$$

$$H * g = Hg = \{h * g : h \in H\} \tag{8.47}$$

to denote the left and right coset generated by g, respectively. If G is abelian, then a particular g will generate the same left and right coset. However, if G is not abelian, this might not be the case.

Exercise 8.3.8 Find all left and right cosets of $H = \{\pm 1, \pm i\}$ in the quaternion group. For each possible $g \in Q$, is $gH = Hg$?

Exercise 8.3.9 From the group in Exercise 8.1.19, let $H = \{0, 1\}$, which is a subgroup. Evaluate $2 \circ H$ and $H \circ 2$.

What makes left and right cosets the same for a particular group element is a question for Section 8.5. Because the results we want to derive here are the same for either left or right cosets, we will look only at right cosets. First, all cosets of a given subgroup have the same cardinality.

Exercise 8.3.10 If G is a group and H is a subgroup of G, then $|H| = |Hg|$ for all $g \in G$.

Regardless of whether G is of finite or infinite order, the number of cosets of H in G might be finite. If so, we call the number of cosets of H the *index* of H in G, and we denote this number by $(G : H)$. Naturally, if G is finite, then so is $(G : H)$, and the following theorem is immediate as an implication of Exercise 8.3.10.

Theorem 8.3.11 (Lagrange). In a finite group, the order of a subgroup divides the order of the group.

Proof. Let G be a group, and let H be a subgroup of G. Since all cosets of H are disjoint and have the same cardinality as H, we have that

$$|G| = (G : H) \times |H| \qquad \square$$

If x is an element of a finite group G, then Lagrange's Theorem says that the order of the subgroup generated by x divides the order of G. Since the cardinality of (x) is the same as the order of the element x (Theorem 8.2.21), we have that $o(x) \mid o(G)$. Beginning here, it is possible to prove several results about elements of a group.

EXERCISE 8.3.12 Suppose $o(G) = n$, and $g \in G$. Then $g^n = e$.[11]

EXERCISE 8.3.13 Suppose G is cyclic of order n and k is a positive integer. If $k \mid n$, then there exists a subgroup of G of order k.

EXERCISE 8.3.14 A group of prime order is cyclic.[12]

8.4 Permutation Groups

In this section, we take an in-depth look at a particular group and two particularly important subgroups that derive from it. Our main purpose is to become familiar with an especially important non-abelian group. Then in Section 8.5, we will use this group as motivation for the definition of a normal subgroup. Requiring a subgroup to be normal is just the right thing to patch up the hole we left in our derivation of quotient groups by requiring that the group be abelian.

8.4.1 Permutation Groups Defined

Let A be a non-empty set, and let S be the set of all bijections from A to itself. We want to take a close look at the group formed on S with the operation of composition. First, note that composition is well defined and closed on S by Exercise 8.1.4, so that S with the operation of composition has properties G1 and G2.

The fact that composition is associative (G3) is a mere exercise in the manipulation of parentheses. For if we pick any $a \in A$, then

$$[(f \circ g) \circ h](a) = (f \circ g)[h(a)] = f(g(h(a))) = f[(g \circ h)(a)] = [f \circ (g \circ h)](a) \tag{8.48}$$

Thus $[(f \circ g) \circ h](a) = [f \circ (g \circ h)](a)$ for all $a \in A$, so that $(f \circ g) \circ h = f \circ (g \circ h)$, and $\circ$ is associative on S. Writing $i : A \to A$, the identity function,

[11] Apply Theorem 8.2.21 and Lagrange's Theorem to (g).

[12] Any element except e is a generator.

then for a given $f \in S$ and any $a \in A$,

$$(f \circ i)(a) = f[i(a)] = f(a) \quad \text{and} \quad (i \circ f)(a) = i[f(a)] = f(a) \tag{8.49}$$

Thus $i \circ f = f \circ i = f$ for all $f \in S$ and i is the identity element (G4). Furthermore, by Theorem 4.4.7, every $f \in S$ has an inverse $f^{-1} \in S$. Since

$$f[f^{-1}(a)] = a = i(a) \quad \text{and} \quad f^{-1}[f(a)] = a = i(a) \tag{8.50}$$

for all $a \in A$, we have $f \circ f^{-1} = f^{-1} \circ f = i$ (G5). Thus $(S, \circ, i, ^{-1})$ is a group, called the *permutation group* on A, and elements of S are called *permutations* of A.

8.4.2 The Symmetric Group

Let $\mathbb{N}_n = \{1, 2, 3, \ldots, n\}$. The permutation group on $\mathbb{N}_n$ is denoted S_n and is called the *symmetric group* on n elements. To be concrete, we will work here with S_6. One way to visualize what a particular $f \in S_6$ does is to imagine six fixed slots, numbered 1–6, with some sort of object in each position. If $f(2) = 5$, this means that the object in slot 2 is moved to slot 5. Thus every element of S_6 does a sort of fruitbasket turnover of the ordered numbers $(1, 2, 3, 4, 5, 6)$. Several notations can describe this effect. One is to display the image of every element of $\mathbb{N}_n$ like this:

$$f = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 \\ 3 & 1 & 2 & 4 & 6 & 5 \end{pmatrix} \tag{8.51}$$

which means $f(1) = 3$, $f(2) = 1$, $f(3) = 2$, $f(4) = 4$, $f(5) = 6$ and $f(6) = 5$. When we compose two permutations in S_6, we work right to left as usual. For example, if we write

$$g = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 \\ 3 & 1 & 4 & 6 & 2 & 5 \end{pmatrix} \tag{8.52}$$

we may calculate $f \circ g$ element by element. To find $(f \circ g)(1)$, we see that 1 is first mapped to 3 by g, and then 3 is mapped to 2 by f. So $(f \circ g)(1) = 2$. Doing the same for the remaining elements yields

$$\begin{aligned} f \circ g &= \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 \\ 3 & 1 & 2 & 4 & 6 & 5 \end{pmatrix} \circ \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 \\ 3 & 1 & 4 & 6 & 2 & 5 \end{pmatrix} \\ &= \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 \\ 2 & 3 & 4 & 5 & 1 & 6 \end{pmatrix} \end{aligned} \tag{8.53}$$

EXERCISE 8.4.1 Calculate $g \circ f$ to show that $(S_6, \circ, i, ^{-1})$ is not abelian.

By flipping the expression for f in Eq. (8.51) upside down and reordering the columns, we see that

$$f^{-1} = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 \\ 3 & 1 & 2 & 4 & 6 & 5 \end{pmatrix}^{-1} = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 \\ 2 & 3 & 1 & 4 & 6 & 5 \end{pmatrix} \tag{8.54}$$

Another way to describe an element of S_6 more succinctly is with *cycle notation*, where $f = (132)(56)$ means 1 is mapped to 3, 3 is mapped to 2, and 2 is mapped to 1 for one cycle. The other cycle indicates 5 is mapped to 6 and 6 is mapped to 5. The absence of 4 means it is mapped to itself. Since (132) and (56) are disjoint, it doesn't matter which you write first. Nor does it matter whether you think of f as a single permutation or as the composition of (132) and (56). If you want, you can scroll the numbers in a cycle around to make it start with any element in the cycle. So $(132) = (213) = (321)$. The identity mapping in S_n is generally written (1). If you compose several permutations in cycle notation that are not disjoint, you can simplify the expression into disjoint cycles.

Example 8.4.2 Simplify $f = (132)(14)(24)(23)(563)$.

Solution Find $f(1)$ first by tracing 1 through the cycles. Since composition of functions is performed from right to left, we see that 1 is first mapped to 4 by (14) and then 4 is not mapped elsewhere by (132). Thus $f(1) = 4$. Now find $f(4)$ by noting that 4 is mapped to 2 by (24), then 2 is mapped to 1 by (132). Thus $f(4) = 1$, and we have completed one cycle (14). Continuing with 2, we see that 2 is mapped to 3 by (23), then 3 is mapped to 2 by (132). Thus $f(2) = 2$, and we can omit it in the cycle notation. Continuing with 3, we see $f(3) = 5$, $f(5) = 6$, and $f(6) = 3$. Thus the simplified expression is $f = (14)(356)$. ■

EXERCISE 8.4.3 Simplify (253)(12)(45)(36)(36) into disjoint cycles.

How many elements are there in S_6? What about in S_n for any n? By Exercise 4.8.9, there are $P(n, n) = n!$ bijections from an n-element set to itself. Another way to conceptualize the question is to ask how many ways there are to arrange the elements of $\mathbb{N}_n$ in the bottom row of Eq. (8.51). There are n possible entries for the first slot, then $n - 1$ remaining possibilities for the second slot, and so on. Multiplying these, we see $|S_n| = n!$. Thus S_n is a group on $n!$ elements, and if $n \geq 3$, S_n is not abelian.

Example 8.4.4 In S_6, determine the subgroup generated by $f = (132)(56)$.

Solution Since S_6 is finite, f has finite order. By Theorem 8.2.20, the subgroup generated by f contains all powers of f.

$$\begin{aligned}
f^2 &= f \circ f = (132)(56) \circ (132)(56) = (123) \\
f^3 &= f^2 \circ f = (123) \circ (132)(56) = (56) \\
f^4 &= f^3 \circ f = (56) \circ (132)(56) = (132) \\
f^5 &= f^4 \circ f = (132) \circ (132)(56) = (123)(56) \\
f^6 &= f^5 \circ f = (123)(56) \circ (132)(56) = (1)
\end{aligned} \tag{8.55}$$

Thus

$$\begin{aligned}
(f) &= \{(1), f, f^2, f^3, f^4, f^5\} \\
&= \{(1), (132)(56), (123), (56), (132), (123)(56)\}
\end{aligned} \tag{8.56}$$

■

Exercise 8.4.5 In S_6, determine the subgroup generated by (125)(346).

Exercise 8.4.6 For $g = (12)$ and $h = (34)$ in S_4, determine $(\{g, h\})$, the subgroup of S_4 generated by g and h.

Exercise 8.4.7 In S_4, let $f = (123)$ and $g = (14)$. Determine whether g generates the same left and right cosets of the subgroup generated by f.

8.4.3 The Alternating Group

Now let's look at an important subgroup of S_n. A permutation of the form (ij) is called a *transposition* because all it does is switch the positions of two elements. Notice $(ij)^{-1} = (ij)$ and $(ij) = (ji)$. If $i = 1$ (or $j = 1$, it doesn't matter), then (ij) becomes $(1j)$. If neither i nor j is 1, then

$$(ij) = (1i)(1j)(1i) \tag{8.57}$$

By Eq. (8.57), every transposition that does not involve 1 can be written as a product of three transpositions, each of which involves 1. So if you're playing some game where objects are lined up in positions $1, \ldots, n$ and you want to swap the positions of the objects in positions i and j, it is possible to do it even if you restrict yourself to swapping an object's position only with the object in position 1. It just takes three moves instead of one. But notice that both (ij) and its equivalent in (8.57) use an odd number of transpositions.

Even though (ijk) is not a transposition, it can be written as a composition of two transpositions. One way to do it is

$$(ijk) = (ik)(ij) \tag{8.58}$$

If none of $\{i, j, k\}$ is 1, what does (8.58) become if we require that all transpositions involve 1, as in (8.57)? Taking a hint from Eq. (8.57) and applying it to both (ik)

and (ij), we can write

$$(ijk) = (ik)(ij) = (1i)(1k)(1i)(1i)(1j)(1i) \tag{8.59}$$

Since $(1i)$ is its own inverse, the two transpositions in the middle of Eq. (8.59) cancel to yield

$$(ijk) = (1i)(1k)(1j)(1i) \tag{8.60}$$

Now $(ijk) = (jki) = (kij)$. So if one of $\{i, j, k\}$ is 1, we may assume $i = 1$, and observe that $(1jk) = (1k)(1j)$. But notice this one fact. Every form of (ijk) that we have written here involves an even number of transpositions.

Suppose we now consider the cycle $\sigma = (x_1, x_2, \dots, x_m)$, where no $x_k = 1$. Can you take a hint from Eq. (8.60) and jump right to a similar form for σ that involves only transpositions of the form $(1x_k)$? How about

$$\sigma = (1x_1)(1x_m)(1x_{m-1})(1x_{m-2}) \cdots (1x_2)(1x_1) \tag{8.61}$$

If some $x_k = 1$, rotate the entries of σ so that $x_1 = 1$. Then Eq. (8.61) works by deleting the transpositions on each end. Be assured there are many other answers. One thing we would like to know is whether all the ways of writing σ with transpositions have something in common.

From all this work, we can see that any element of S_n, once written in cycle notation with disjoint cycles, can be broken down into a composition of transpositions in at least some way. The identity can be thought of as requiring zero transpositions, or if you prefer, we can write $(1) = (12)(12)$ for $n \geq 2$. One characteristic of elements of S_n that we want to point out, but not prove in this text, is that of all possible decompositions of a particular permutation into transpositions, either they will all involve an even number of transpositions or they will all involve an odd number. This is not a trivial fact to demonstrate. So we will just accept it here and leave the proof to your later coursework in algebra. Thus the $n!$ elements of S_n are partitioned into two classes. If f always decomposes into an even number of transpositions, then f is called an *even permutation*. Similarly, if f always decomposes into an odd number of transpositions, then f is called an *odd permutation*.

Now let's create an important subset of S_n. Let

$$A_n = \{f \in S_n : f \text{ is an even permutation}\} \tag{8.62}$$

What do you think the relationship between A_n and S_n is?

EXERCISE 8.4.8 A_n is a subgroup of S_n.

We call A_n the *alternating group* on n elements.

Given that $|S_n| = n!$, how many elements do you think A_n has? Your first thought might be that, since every element of S_n is either even or odd, it would be

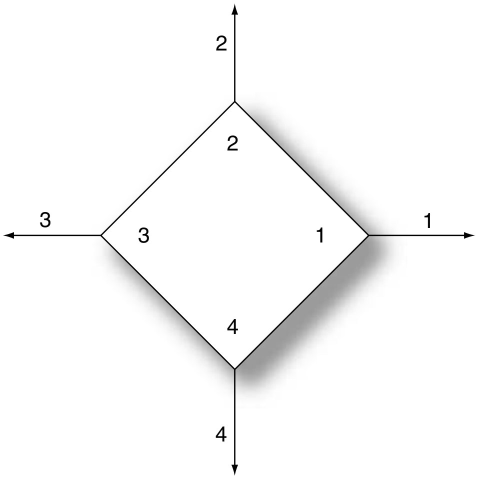

Figure 8.1 Rigid square used to generate the dihedral group.

only natural that half would be even and half would be odd, so that $|A_n| = n!/2$. If so, you're right. In Exercise 8.3.10, you showed that all cosets of a subgroup have the same cardinality. So if you can show that A_n has precisely two cosets, itself and its complement, it will follow that $|A_n| = n!/2$.

EXERCISE 8.4.9 Show that $|A_n| = n!/2$ by showing that A_n has precisely two cosets.[13]

8.4.4 The Dihedral Group

Now let's look at another important subgroup, this time for the particular group S_4. Consider a rigid square with the numbers $1, 2, 3, 4$ etched on its corners, sitting in the xy-plane where the positions of each corner are also written in the plane. (See Fig. 8.1.) Consider the following two moves for the square.

1. A rotation 90° counterclockwise, a move we will call ρ (the Greek letter rho).
2. A flip upside down, sending the top corner to the bottom, and vice versa, and leaving the side corners fixed. Call this move ϕ (the Greek letter phi).

The move ρ can be expressed as an element of S_4: $\rho = (1\,2\,3\,4)$. Visualize ρ and the cycle notation expression of it as sending the corner that initially occupies position 1 in the xy-plane to position 2, the corner initially in position 2 to position 3, and so on. Similarly, $\phi = (2\,4)$, which sends the corner initially in position 2 to position 4, and vice versa. If we consider all possible ultimate positions of the

[13] When do f_1 and f_2 generate the same coset of A_n? See the comments after Definition 8.3.5.

square that can result from combining moves ρ and ϕ in any way by composition, we have created a way of visualizing the subgroup of S_4 generated by $\{\rho, \phi\}$. This subgroup of S_4 is called the *dihedral group* and is denoted D_8.

First, let's note how many elements there are in D_8. We ask how many possible ultimate positions are there for the square that can result from a combination of rotations and flips. Perhaps it is clear that the square can end up either top side up or top side down, and in any one of four states of rotation. Thus eight distinguishable ultimate positions are possible, and $o(D_8) = 8$.

Let's create a Cayley table for D_8 in the following way. Each element of D_8, that is, each ultimate position for the square, can be obtained by doing any necessary rotation first, then flipping afterward, if necessary. Thus every element of D_8 can be written in the form $\phi^m \rho^n$. As an example, $\phi\rho^3$ would rotate 270° counterclockwise and then flip. This maneuver could be written in cycle notation as $(24)(1234)^3$.

EXERCISE 8.4.10 Simplify $\phi\rho^3$ into cycle notation with disjoint cycles.

Since ρ has order 4 and ϕ has order 2, every element of D_8 can be written as $\phi^m \rho^n$, where $0 \le m \le 1$ and $0 \le n \le 3$. Writing elements of D_8 in this way and denoting the identity by i, we have

$$D_8 = \{i, \rho, \rho^2, \rho^3, \phi, \phi\rho, \phi\rho^2, \phi\rho^3\} \tag{8.63}$$

To fill in the Cayley table values, we must combine all elements of D_8 and write the compositions in a form from Eq. (8.63). This takes a little bit of work but yields some useful principles as you work through it. For example, what is $(\phi\rho) \circ (\phi) = \phi\rho\phi$? The first maneuver is a flip, followed by a 90° rotation and flip.

EXERCISE 8.4.11 Which element of D_8 from (8.63) is equivalent to $\phi\rho\phi$?

Table 8.64 contains a few of the entries for the Cayley table of D_8. In the next exercise, you will finish out this table and be pointed to a systematic approach that can save you a lot of time. Remember! In a Cayley table, the entry down the left column is written on the left, and the entry across the top row is written on the right. Since composition is read from right to left, it means that the entry from the top row is actually performed first!

$\circ$	i	ρ	ρ^2	ρ^3	ϕ	$\phi\rho$	$\phi\rho^2$	$\phi\rho^3$
i	i	ρ	ρ^2	ρ^3	ϕ	$\phi\rho$	$\phi\rho^2$	$\phi\rho^3$
ρ	ρ	ρ^2	ρ^3	i	$\phi\rho^3$	ϕ		
ρ^2	ρ^2	ρ^3	i	ρ	$\phi\rho^2$			
ρ^3	ρ^3	i	ρ	ρ^2				
ϕ	ϕ							
$\phi\rho$	$\phi\rho$							
$\phi\rho^2$	$\phi\rho^2$							
$\phi\rho^3$	$\phi\rho^3$							

(8.64)

EXERCISE 8.4.12 Complete Table 8.64.[14]

8.5 Normal Subgroups

Let's return to quotient groups and recall from Section 8.3 how our assumption that G is abelian got us over the hump of showing that the operation on G/H from Definition 8.3.5 is well defined. For notational simplicity, we write

$$Ha * Hb = H(ab) \tag{8.65}$$

using juxtaposition for the binary operation on G and in the coset notation. To show $*$ is well defined on G/H, we would suppose

$$Ha_1 = Ha_2 \quad \text{and} \quad Hb_1 = Hb_2 \tag{8.66}$$

and try to show from this that

$$H(a_1b_1) = H(a_2b_2) \tag{8.67}$$

Since Eq. (8.67) is a set equality, one way to approach our task is to chase an element from one side to the other and back. A more efficient approach, however, is to exploit Exercise 8.3.3 and the fact that equivalence classes partition a set. In other words, to show that Eqs. (8.66) imply Eq. (8.67), it suffices to show that if $a_1 \equiv_H a_2$ and $b_1 \equiv_H b_2$, then $a_1b_1 \equiv_H a_2b_2$. Let's try to do that now, and see where we get stuck in the absence of G being abelian.

Suppose $a_1 \equiv_H a_2$ and $b_1 \equiv_H b_2$. Then $a_1a_2^{-1}, b_1b_2^{-1} \in H$. We need to show that $(a_1b_1)(a_2b_2)^{-1} \in H$, which is equivalent to $a_1b_1b_2^{-1}a_2^{-1} \in H$. Perhaps the convenience of commutativity is evident at this point because the only assumptions we have involve $a_1a_2^{-1}$ and $b_1b_2^{-1}$. We must proceed with extreme caution, so we should spend some time thinking out loud.

First, we know that there exists $h_1 \in H$ such that $b_1b_2^{-1} = h_1$. Thus our task now is to show that $a_1h_1a_2^{-1} \in H$. At this point we're stuck because the only way to apply our assumptions to an expression involving a_1 and a_2^{-1} is if they can be associated. In our case, h_1 is in the way. One possible assumption that would get us past this problem would be to assume that h_1 commutes with a_2^{-1}, which would require us to assume that every element of H commutes with every element of G. This is a pretty strong assumption about H, and we can do better.

The standard assumption about H, weak as possible but still strong enough to give us a well-defined operation on G/H, is to assume a way to reorder terms

[14] First convert all expressions of the form $\rho^i\phi$ to their equivalents in the form $\phi\rho^j$. Use these results to take expressions of the form $\phi^s\rho^t\phi^u\rho^v$ and reorder $\rho^t\phi^u$ in the middle. Thus all ρs will gravitate to the right and all ϕs to the left.

where we can replace $h_1a_2^{-1}$ with $a_2^{-1}h_2$ for some $h_2 \in H$. If we have this, then $a_1h_1a_2^{-1} = a_1a_2^{-1}h_2$, and the fact that $a_1a_2^{-1} \in H$ allows us to conclude that

$$(a_1b_1)(a_2b_2)^{-1} = a_1b_1b_2^{-1}a_2^{-1} = a_1h_1a_2^{-1} = a_1a_2^{-1}h_2 \in H \tag{8.68}$$

So here is the demand on H that will get us over the hump of showing that the operation on G/H is well defined when G is not abelian. Since a_2^{-1} in our analysis above could be any element of the group, and h_1 could be any element of H, we arrive at the following definition of a special sort of subgroup where we can be sure that the operation on G/H is well defined.

Definition 8.5.1 Suppose H is a subgroup of a group G. If for all $g \in G$ and $h \in H$, there exists $h_1 \in H$ such that $hg = gh_1$, then H is called a *normal subgroup* of G, and we write $H \lhd G$.

Definition 8.5.1 only allows you to swap hg for gh_1. It does not explicitly allow you to swap gh with some h_1g. But you can show that it does.

Exercise 8.5.2 Suppose H is a normal subgroup of G. Then for all $g \in G$ and $h \in H$, there exists $h_1 \in H$ such that $gh = h_1g$.[15]

All our work here shows that the operation defined in Eq. (8.65) is well defined if H is a normal subgroup of G. Furthermore, your other work from Exercise 8.3.4 where you did not exploit the abelian nature of G completes the proof that G/H is a group, and we arrive at the following.

Theorem 8.5.3 Suppose G is a group and H is a normal subgroup of G. Then G/H with binary operation $*$ defined by $(Ha) * (Hb) = H(ab)$ is a group with identity He and inverses $(Ha)^{-1} = Ha^{-1}$.

There are ways other than Definition 8.5.1 to define normal subgroup, and different authors approach the idea in different ways. As we will see in Section 8.6, normality can be naturally defined in terms of mappings between groups. But we have chosen our definition, so any equivalent forms will have to be demonstrated as theorems. If we reread the analysis that led us to Definition 8.5.1, another way to state what we needed might jump out at us. Rather than state that there exists $h_2 \in H$ such that $a_2^{-1}h_2 = h_1a_2^{-1}$, we simply could have insisted that there exists $h_2 \in H$ such that $h_2 = a_2h_1a_2^{-1}$, or simply that $a_2h_1a_2^{-1} \in H$. This has a natural appeal because testing a subgroup for normality then reduces simply to making sure that the criterion in the following theorem is satisfied.

[15] Apply Definition 8.5.1 to $h^{-1}g^{-1}$.

Theorem 8.5.4 Suppose H is a subgroup of G. Then H is normal if and only if for all $h \in H$ and $g \in G$, $g^{-1}hg \in H$.

It is one thing to say that a subgroup is closed under the operations of the group. But the fact that $g^{-1}hg \in H$ for all $h \in H$ and $g \in G$ adds a new dimension to the closure of H by saying, in a sense, that elements of H cannot be kicked outside of H by boxing them inside things of the form $g^{-1}\square g$. For a given $h \in H$, an expression of the form $g^{-1}hg$ is called a *conjugate* of h. Theorem 8.5.4 says that H is normal if and only if the conjugates of all its elements lie in H. Thus choosing any $h \in H$ and $g \in G$, it might be that $g^{-1}hg$ is different from h, but at least it's still in H. This is the way you will want to argue the following.

Exercise 8.5.5 The alternating group A_n is a normal subgroup of S_n.

If G is not abelian and a subgroup H is not normal, then a conjugate of an element of H might or might not be in H. For example, using $\rho = (1234) \in D_8$ and $(12) \in S_4$,

$$(12)^{-1}(1234)(12) = (12)(1234)(12) = (1342) \notin D_8 \tag{8.69}$$

which by Theorem 8.5.4 proves that D_8 is not normal as a subgroup of S_4. On the other hand, using $(123) = (13)(12) \in A_4$ and $(14) \in S_4$,

$$(14)^{-1}(123)(14) = (14)(123)(14) = (234) = (24)(23) \in A_4 \tag{8.70}$$

which is a consequence of the fact that A_4 is a normal subgroup of S_4.

Exercise 8.5.6 Use Theorem 8.5.4 to show that $H = \{(1), (123), (132)\}$ is a normal subgroup of S_3.

Exercise 8.5.7 Is $H = \{(1), (12)\}$ a normal subgroup of S_3? Prove or disprove.

Exercise 8.5.8 From Exercise 8.4.6, $H = \{(1), (12), (34), (12)(34)\}$ is a subgroup of S_4. Prove or disprove that H is normal in S_4.

If G is abelian, conjugation is not particularly interesting.

Theorem 8.5.9 A group G is abelian if and only if every $h \in G$ has no conjugates other than itself.

The proof of Theorem 8.5.9 should be immediately clear, for the conditions that $gh = hg$ for all $g, h \in G$ and $g^{-1}hg = h$ for all $g, h \in G$ are identical.

If we think of g as fixed and allow h to take on all values in H, we create what we call a conjugate of the subgroup H. Notationally, we write this as

$$g^{-1}Hg = \{g^{-1}hg : h \in H\} \tag{8.71}$$

One interesting characteristic of the conjugate of a subgroup is that it is also a subgroup of G.

EXERCISE 8.5.10 Let H be a subgroup of a group G and fix $g \in G$. Then $g^{-1}Hg$ is a subgroup of G.

Given two subgroups H_1 and H_2, to say that H_2 is a conjugate of H_1 therefore means that there exists some $g \in G$ such that $H_2 = g^{-1}H_1g$.

EXERCISE 8.5.11 Let G be a group, and let H_1 and H_2 be subgroups of G. Define $H_2 \equiv H_1$, provided H_2 is a conjugate of H_1. Show that $\equiv$ is an equivalence relation on the family of all subgroups of G.

EXERCISE 8.5.12 For $H = \{(1), (12), (34), (12)(34)\}$, which is a subgroup of S_4, determine the conjugate subgroup $(13)^{-1}H(13)$.

For another example of a conjugate subgroup, consider D_8 as a subgroup of S_4 and let $g = (12)$. For a given $\delta \in D_8$, the expression $g^{-1}\delta g$ can be thought of in the following way. Since g acts first, it switches the numbers in positions 1 and 2 on the square (an illegal move in D_8). Then δ does a rotation and/or flip on this newly labeled square. Finally, g^{-1} switches again the numbers in positions 1 and 2 on the square. For example,

$$g^{-1}\rho g = (12)(1234)(12) = (1342) \notin D_8 \tag{8.72}$$

Transforming all the elements of D_8 in this way creates another subgroup of S_4 that you might need to play around with to understand. It turns out that $g^{-1}D_8g$ as an algebraic structure is just like D_8, except that the rigid square we described before will not work as a way to visualize it. Instead, picture the numbers $\{1, 2, 3, 4\}$ being pushed from corner to corner by $g^{-1}\rho g$ and $g^{-1}\phi g$ according to Figure 8.2. The

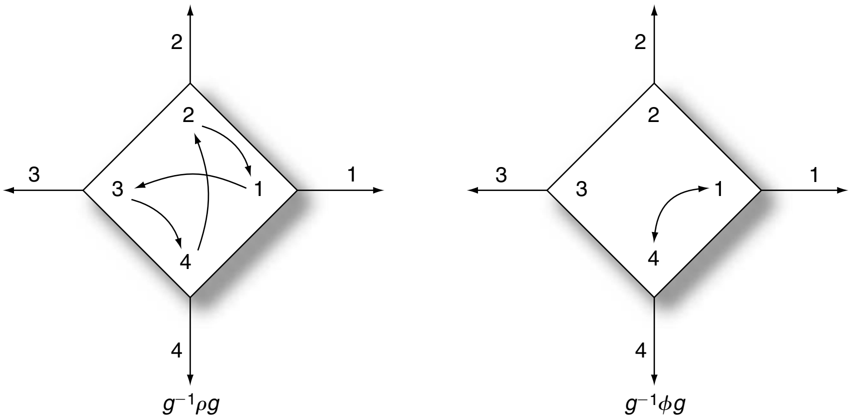

Figure 8.2 Effects of $g^{-1}\rho g$ and $g^{-1}\phi g$ on the square.

point to be made is that $g^{-1}D_8g$ is different from D_8, though it is a subgroup of S_4. They both contain the identity, but except for that overlap, they cut through S_4 in completely different directions.

If H is a normal subgroup of G, the situation is a little different concerning conjugates of H. The proof of the following should be quick. It says H is normal in G if and only if H has no conjugates other than itself.

Exercise 8.5.13 Suppose H is a subgroup of G. Then H is normal if and only if $g^{-1}Hg = H$ for all $g \in G$.

In showing that the binary operation on a quotient group is well defined, another feature of H that would get us over the hump from h_2a_2 to a_2h_3 involves linking the right coset Ha_2 to the left coset a_2H. That is, if we had been given that $Ha_2 = a_2H$, our problem would have been solved in precisely the same way. So here is another equivalence of normality.

Exercise 8.5.14 Suppose H is a subgroup of G. Then H is normal if and only if $Hg = gH$ for all $g \in G$.

To sum up, notice how our retreat from the requirement that G be abelian motivated the definition of a characteristic of subgroups that allows us still to construct the quotient group. And note the following.

Exercise 8.5.15 Every subgroup of an abelian group is normal.

Since our definition of normal subgroup was inspired by a retreat from the global condition of G being abelian, let's compare and contrast normality in all its forms to similarly worded statements about abelian groups.

	If G is abelian and $H < G$:	If $H \lhd G$:
1.	For all $g, h \in G$, $gh = hg$.	For all $g \in G$ and $h \in H$, there exists $h_1 \in H$ such that $hg = gh_1$.
2.	For all $g, h \in G$, $g^{-1}hg = h$.	For all $g \in G$ and $h \in H$, $g^{-1}hg \in H$.
3.	For all $g \in G$, $g^{-1}Hg = H$.	For all $g \in G$, $g^{-1}Hg = H$.
4.	For all $g \in G$, $Hg = gH$.	For all $g \in G$, $Hg = gH$.

Exercise 8.5.16 In Exercise 8.2.8, you showed that the center of a group G is a subgroup of G. Prove that the center of G is normal in G.

Exercise 8.5.17 In Exercise 8.2.9, you showed that the intersection across a family of subgroups of G is itself a subgroup of G. Show that the intersection across a family of normal subgroups of G is normal in G.

As a final observation, consider four sets $H \subseteq H_1 \subseteq G \subseteq G_1$, where H is a normal subgroup of G. Normality of H as a subgroup of G says something about the way elements of H behave in the presence of elements of G, and not simply how they behave among themselves. So, if H_1 is also a group, then the fact that H is normal in G will imply H is normal in H_1 as well. For if $g^{-1}hg \in H$ for all $g \in G$, then certainly the same is true for all $g \in H_1$. Now if G_1 is a group, we know that H is a subgroup of G_1. However, it might be that H is not normal in G_1. Even though $g^{-1}hg \in H$ for all $g \in G$, there might be some $g \in G_1 - G$ for which $g^{-1}hg \notin H$.

Example 8.5.18 Beginning with S_4, create the following subgroups. First, let $H = \{(1), (12)\}$. Since (12) is its own inverse, H is a subgroup of S_4. Let $G = \{(1), (12), (34), (12)(34)\}$. By Exercise 8.4.6, G is a subgroup of S_4. Furthermore, G is abelian. By Exercise 8.5.15, H is a normal subgroup of G because G is abelian. However, by the next exercise, H is not normal in S_4. ■

EXERCISE 8.5.19 Show that $H = \{(1), (12)\}$ is not normal in S_4 by finding a conjugate of (12) that is not in H.

8.6 Group Morphisms

Now that we have a basic understanding of groups and some of their internal structure, let's turn our attention outward to a special type of function from one group to another. The special feature we want these functions to have is that they preserve the binary operation.

Definition 8.6.1 Suppose $(G, *, e_G, ^{-1})$ and $(H, \cdot, e_H, ^{-1})$ are groups, and suppose $\phi : G \to H$ is a function with the property that $\phi(x * y) = \phi(x)\cdot\phi(y)$ for all $x, y \in G$. Then ϕ is called a *homomorphism*, or simply a *morphism* from G to H. If ϕ is one-to-one, it is called a *monomorphism*. If ϕ is onto, it is called an *epimorphism*. If ϕ is both one-to-one and onto, it is called an *isomorphism*, and we write $G \cong H$, which is read "G is isomorphic to H." If $\phi : G \to G$ is an isomorphism, ϕ is called an *automorphism*.

EXERCISE 8.6.2 Let $\mathbb{Z}$ be the group of integers under addition and E the group of all even integers under addition. Assuming that each of the following is a function, prove the following.

(a) Show that $\phi_1 : \mathbb{Z} \to E$ defined by $\phi_1(n) = 2n$ is an isomorphism.

(b) Show that $\phi_2 : \mathbb{Z} \to E$ defined by $\phi_2(n) = 4n$ is a monomorphism that is not an isomorphism.

(c) Show that $\phi_3 : \mathbb{Z} \to \mathbb{Z}$ defined by $\phi_3(n) = -n$ is an automorphism.

(d) Show that $\phi_4 : \mathbb{Z} \to \mathbb{Z}$ defined by $\phi_4(n) = 2n + 2$ is not a morphism.

Exercise 8.6.3 Let $\mathbb{Z}$ be the group of integers under addition, let n be a positive integer ($n \geq 2$). Show that $\phi : \mathbb{Z} \to \mathbb{Z}_n$ defined by $\phi(k) = (n) + k$ is an epimorphism that is not an isomorphism.

Example 8.6.4 Let $G = (\mathbb{R}^+, \times, 1, ^{-1})$ and $H = (\mathbb{R}, +, 0, -)$. Even though we have not discussed logarithms in this text, your work in pre-calculus reveals that $\phi(x) = \ln x$ is an isomorphism. For $\ln x$ is a bijection from $\mathbb{R}^+$ to $\mathbb{R}$ and satisfies $\phi(xy) = \ln(xy) = \ln x + \ln y = \phi(x) + \phi(y)$. ■

Example 8.6.5 Let G and H be groups, and define $\phi : G \to H$ by $\phi(x) = e_H$ for all $x \in G$. Then ϕ is called the *trivial* morphism. ■

Exercise 8.6.6 Define addition in $\mathbb{R}^2$ by

$$(x_1, y_1) + (x_2, y_2) = (x_1 + x_2, y_1 + y_2)$$

which is clearly a group. Let $\phi : \mathbb{R}^2 \to \mathbb{R}$ be defined by $\phi[(x, y)] = x$. Show that ϕ is an epimorphism that is not an isomorphism from the additive group $\mathbb{R}^2$ to the additive group of real numbers. (This is called a *projection morphism* because every point in the xy-plane is mapped directly up or down to its x-coordinate.)

Exercise 8.6.7 Let G be a group, and define $\phi : G \to G$ by $\phi(g) = g^{-1}$. Assuming ϕ is a bijection from G to itself, show that ϕ is an automorphism of G if and only if G is abelian.

The morphic behavior of $\phi : G \to H$ does not explicitly require that particular elements of G must map to particular elements of H. However, there are some restrictions of this sort inherent in the definition.

Theorem 8.6.8 If $\phi : G \to H$ is a group morphism, then:

1. $\phi(e_G) = e_H$.
2. For all $x \in G$, $\phi(x^{-1}) = [\phi(x)]^{-1}$.

We will prove part (1) and leave part 2 to you as an exercise.

Proof of Part (1).

$$e_H \cdot \phi(e_G) = \phi(e_G) = \phi(e_G * e_G) = \phi(e_G) \cdot \phi(e_G) \tag{8.73}$$

By cancellation, $\phi(e_G) = e_H$. □

Exercise 8.6.9 Prove part (2) of Theorem 8.6.8.

By using part (1) of Theorem 8.6.8 as the root of an induction argument for $n \geq 0$, then using part (2) to take care of $n < 0$, you can show the following.

EXERCISE 8.6.10 If $\phi : G \to H$ is a group morphism and $g \in G$, then for all integers n,

$$\phi(g^n) = [\phi(g)]^n \tag{8.74}$$

EXERCISE 8.6.11 Restate Theorem 8.6.8 and Exercise 8.6.10 in their additive forms.

The statement $G \cong H$ in Definition 8.6.1 looks like a form of equivalence. It is a statement about the existence of a bijection from G to H, with the additional property that it preserves the binary operation. By Exercise 4.4.9, most of the work in proving that $\cong$ has properties E1–E3 has been done and will therefore require only references to applicable theorems. However, in each property E1–E3, something will have to be shown about the morphic behavior of the functions involved.

EXERCISE 8.6.12 The relation $\cong$ in Definition 8.6.1 is an equivalence relation on the set of all groups.

The word *isomorphism* has a connotation to it that deserves pointing out. For G to be isomorphic to H means that G and H are essentially the same group, in the sense that all the elements from one can be swapped one for one with those in the other and the internal relationships between them as expressed by $*$ are retained by $\cdot$. By giving a new name to every $x \in G$, namely, $\phi(x)$, and swapping the binary operation symbol in G for the symbol in H, we have effectively dressed $(G, *, e_G, ^{-1})$ in the clothing of $(H, \cdot, e_H, ^{-1})$. So to be able to map every element of G to a unique and distinct element of H, exhausting all elements of H and preserving the binary operation, is to show that, as far as their structure as groups is concerned, they are identical. Here are some illustrations of this principle.

EXERCISE 8.6.13 Suppose $\phi : G \to H$ is a group isomorphism. Then,

(a) If G is abelian, then H is abelian.

(b) If G is cyclic, then H is cyclic.

Instead of writing $\text{Rng}(\phi)$, we will usually write $\phi(G)$. Subgroups of the domain and codomain are related in the following.

EXERCISE 8.6.14 Suppose $\phi : G \to H$ is a group morphism. Then

(a) $\phi(G)$ is a subgroup of H.

(b) If N is a normal subgroup of G, then $\phi(N)$ is a normal subgroup of $\phi(G)$.

(c) If N is a normal subgroup of H, then $\phi^{-1}(N)$ is a normal subgroup of G.

Notice that part (b) of Exercise 8.6.14 does not say that $\phi(N)$ is a normal subgroup of H. If ϕ is not onto, it is possible that $\phi(N)$ is normal in $\phi(G)$ but not normal in H.

Theorem 8.6.8 ensures that e_G always maps to e_H under a group morphism. It might be that other elements of G also map to the identity in H. In Exercise 8.6.3, every multiple of n maps to the identity coset $(n) + 0$. In Example 8.6.5, every element of G maps to the identity in H. We give a name to the set of all elements in G that map to the identity by a morphism.

Definition 8.6.15 Suppose $\phi : G \to H$ is a group morphism. Then the pre-image of e_H is called the *kernel* of ϕ and is denoted $\mathrm{Ker}(\phi)$. If $\mathrm{Ker}(\phi)$ contains only e_G, we say that the kernel of ϕ is *trivial.*

Exercise 8.6.16 If $\phi : G \to H$ is a group morphism, then $\mathrm{Ker}(\phi)$ is a normal subgroup of G.

An interesting property of group morphisms is that you can sometimes learn a lot about the behavior of ϕ across all of G by looking at its behavior at certain places in G. Clearly, if ϕ is one-to-one, then $\mathrm{Ker}(\phi)$ is trivial, because then only e_G will map to e_H. But, interestingly, the converse of this is also true.

Exercise 8.6.17 Suppose $\phi : G \to H$ is a group morphism. If $\mathrm{Ker}(\phi)$ is trivial, then ϕ is one-to-one.

So if $\phi^{-1}(e_H)$ has only one element, then $\phi^{-1}(y)$ has only one element for every $y \in \phi(G)$. We can go even further to show that

$$\left|\phi^{-1}(y)\right| = |\mathrm{Ker}(\phi)| \tag{8.75}$$

for all $y \in \phi(G)$, so that all pre-image sets of individual elements in the range of ϕ have the same cardinality. Probably the easiest way to do this is to exploit the fact that the kernel is a normal subgroup of G in order to prove something even stronger than Eq. (8.75). The next exercise says that the pre-image of each element in the range of ϕ is simply one of the cosets of the kernel. By Exercise 8.3.10, all cosets of a subgroup have the same cardinality, so that Eq. (8.75) follows. In the next exercise, we use left cosets to keep the notation uncluttered.

Exercise 8.6.18 If $\phi : G \to H$ is a group morphism, then for all $y \in \phi(G)$, there exists $a \in G$ such that $\phi^{-1}(y) = a\,\mathrm{Ker}(\phi)$.

If you've caught on to what's happening here, you might have observed that any group morphism $\phi : G \to H$ makes some interesting statements about the internal structure of G. A morphism ϕ gives rise to a normal subgroup of G, namely, $\mathrm{Ker}(\phi)$. From there all our theory from Section 8.5 applies to present us with a quotient group $G/\mathrm{Ker}(\phi)$, with its binary operation $a\,\mathrm{Ker}(\phi) * b\,\mathrm{Ker}(\phi) =$

$(ab)\,\mathrm{Ker}(\phi)$, identity $\mathrm{Ker}(\phi)$ and inverses $[a\,\mathrm{Ker}(\phi)]^{-1} = a^{-1}\,\mathrm{Ker}(\phi)$. So any time you are given a group G and can manage to find a morphism into some other group, then you have managed also to find a normal subgroup of G and motivate a quotient group from it.

Also, if you're like a lot of people at your stage of the mathematical game, you wish there were a better way to see what's going on inside this new quotient group $G/\,\mathrm{Ker}(\phi)$ than by visualizing cosets whacking each other around. Well, there is another way to visualize $G/\,\mathrm{Ker}(\phi)$, for it is essentially the same as (isomorphic to) $\phi(G)$. We will present a full-blown statement and proof of that soon.

A group morphism $\phi : G \to H$ gives rise to a normal subgroup of G and a quotient group from it. Now let's go the other way. Given a group G and any normal subgroup N, we can *create* a group H and an epimorphism $\phi : G \to H$ such that $\mathrm{Ker}(\phi) = N$. Showing this does not involve much we have not seen before, mostly just some observations. You will provide the only missing detail as an exercise.

Theorem 8.6.19 If $(G, *, e, ^{-1})$ is a group and N is a normal subgroup of G, then $\phi : G \to G/N$ defined by $\phi(x) = Nx$ is an epimorphism whose kernel is N.

Proof. From our previous work, ϕ is defined on all of G and is well defined. Also, ϕ is onto because every element of G/N is a coset of N, generated by some $x \in G$, and $\phi(x) = Nx$. We must show that ϕ behaves morphically and satisfies $\mathrm{Ker}(\phi) = N$. Pick $x, y \in G$. Since the binary operation on G/N is well defined,

$$\phi(xy) = N * (xy) = Nx * Ny = \phi(x)\phi(y) \tag{8.76}$$

Finally, by Exercise 8.6.20, $\mathrm{Ker}(\phi) = N$. □

EXERCISE 8.6.20 Finish the proof of Theorem 8.6.19 by showing that $\mathrm{Ker}(\phi) = N$.

So we can have it both ways. Two groups and a morphism give rise to a normal subgroup of the domain. And a group with a normal subgroup gives rise to another group and a morphism onto it. Here is the final tie. Whichever you start with and use to create the other, the range of the morphism and the quotient group of G are isomorphic. That is, $\phi(G)$ and $G/\,\mathrm{Ker}(\phi)$ are essentially the same group, as the following theorem states. For notational simplicity, we assume that the given morphism is onto so that we don't have to distinguish between H and $\phi(G)$.

Theorem 8.6.21 Suppose $\phi : G \to H$ is a group epimorphism. Then

$$G/\,\mathrm{Ker}(\phi) \cong H \tag{8.77}$$

We will supply a skeleton of the proof here, though most of the details are left to you in the exercise that follows. As we do so, look at Figure. 8.3, which illustrates all the groups and mappings involved. First, there are the given groups G and H

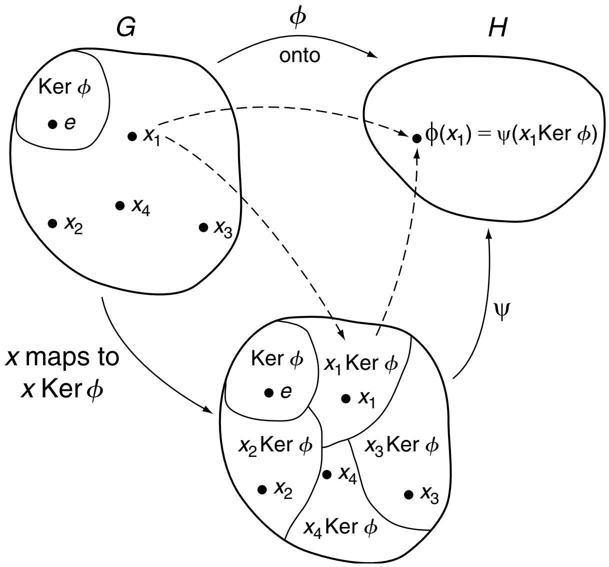

Figure 8.3 Isomorphism between $G/\operatorname{Ker}(\phi)$ and H.

with the epimorphism $\phi : G \to H$. Now create a carbon copy of G, only collect all the elements of $\operatorname{Ker}(\phi)$ and hogtie them together into a single entity in the sketch of $G/\operatorname{Ker}(\phi)$. Similarly, go to each coset of $\operatorname{Ker}(\phi)$, take all its elements, and lump them together into a single entity in $G/\operatorname{Ker}(\phi)$ to create a visualization of $G/\operatorname{Ker}(\phi)$. We have a link between G and $G/\operatorname{Ker}(\phi)$, and that is the function that maps $x \in G$ to $x\operatorname{Ker}(\phi)$. It might be that ϕ is one-to-one, or maybe it is not. But the extent to which ϕ collapses elements of G down to single elements of H is precisely the extent to which elements of G clump together into cosets in $G/\operatorname{Ker}(\phi)$ (Exercise 8.6.18), which itself depends on the size of the kernel. The task is to show that $G/\operatorname{Ker}(\phi)$ and H are isomorphic by finding the required mapping between them. We must map each coset in $G/\operatorname{Ker}(\phi)$ to an element of H, and the way to do this is to choose a coset $x\operatorname{Ker}(\phi)$, grab some element of it, say x, and send the whole coset to $\phi(x)$ in H. In the following proof, we use $\cdot$ and $*_K$ to represent the binary operations in H and $G/\operatorname{Ker}(\phi)$, respectively.

Proof. We must find a bijection $\psi : G/\operatorname{Ker}(\phi) \to H$ such that

$$\psi[x\operatorname{Ker}(\phi) *_K y\operatorname{Ker}(\phi)] = \psi[x\operatorname{Ker}(\phi)] \cdot \psi[y\operatorname{Ker}(\phi)] \qquad (8.78)$$

for all $x\operatorname{Ker}(\phi), y\operatorname{Ker}(\phi) \in G/\operatorname{Ker}(\phi)$. So define $\psi : G/\operatorname{Ker}(\phi) \to H$ by

$$\psi[x\operatorname{Ker}(\phi)] = \phi(x) \qquad (8.79)$$

By Exercise 8.6.22, ψ is well defined on all of $G/\operatorname{Ker}(\phi)$, is one-to-one and onto, and satisfies Eq. (8.78). □

EXERCISE 8.6.22 Finish the proof of Theorem 8.6.21 by showing that $\psi : G/\operatorname{Ker}(\phi) \to H$ defined by Eq. (8.79) satisfies the following.

(a) ψ is defined on all of $G/\operatorname{Ker}(\phi)$ (F1).

(b) ψ is well defined (F2).

(c) ψ is one-to-one.

(d) ψ is onto.

(e) For all $x\operatorname{Ker}(\phi), y\operatorname{Ker}(\phi) \in G/\operatorname{Ker}(\phi)$,

$$\psi[x\operatorname{Ker}(\phi) * y\operatorname{Ker}(\phi)] = \psi[x\operatorname{Ker}(\phi)]\cdot\psi[y\operatorname{Ker}(\phi)]$$

As a final note, the theorems from this section provide us with another logical equivalence to normality of a subgroup. If N is a normal subgroup of G, then there exists a group H and an epimorphism $\phi : G \to H$ such that $\operatorname{Ker}(\phi) = N$. Also, if $\phi : G \to H$ is an epimorphism, $\operatorname{Ker}(\phi)$ is a normal subgroup of G. So given a group G and a subgroup N, N is normal in G if and only if there exists some group H and some morphism $\phi : G \to H$ such that $N = \operatorname{Ker}(\phi)$.

EXERCISE 8.6.23 Construct notation and Cayley tables to determine (up to isomorphism) all groups of five or fewer elements.

EXERCISE 8.6.24 Find isomorphisms between the four-element groups you found in Exercise 8.6.23 and the following.

(a) The multiplicative group $S = \{\pm 1 \pm i\}$

(b) The subgroup of S_4 in Exercise 8.4.6

EXERCISE 8.6.25 Describe all cyclic groups.

EXERCISE 8.6.26 Both the dihedral group (Section 8.4) and the quaternion group (Example 8.1.18) have eight elements. Show that they are not isomorphic.

Rings

We can create algebraic structures of greater complexity than a group by endowing a set with two binary operations and laying down some assumptions about how these operations behave, both on their own and in relation to each other. In this chapter we look at several such structures. Before we do, some explanation is in order about how we will proceed, for the theory of rings involves so many details that a road map will be very helpful.

First, in Section 9.1, we will define the most general algebraic structure with two binary operations, a *ring*, and construct several important examples. At the same time, we will make a passing reference to *fields*, the most specialized kind of ring we will study. In Section 9.2, we define *subring* and construct a number of examples. In Section 9.3, we will look at several properties that the most general rings share. One particularly important class of rings can be created by *adjoining* an element to a given ring; we devote Section 9.4 to this class of examples. In Section 9.5, we dive down inside a ring to look at specialized substructures of a general ring. Ideals, principal ideals, prime ideals, and maximal ideals are special types of substructures we will see there. In Sections 9.8–9.11, we will study four increasingly specialized kinds of rings: integral domains, unique factorization domains, principal ideal domains, and Euclidean domains. Each class of these structures is a proper subset of the class that comes before it, so as we progress, we will demonstrate (or at least refer to) examples that illustrate this. For example, we will see a ring that is not an integral domain, an integral domain that is not a unique factorization domain, and so on. In Section 9.14 we will look at ring morphisms, and finally, in Section 9.15, we will build quotient rings.

9.1 Rings and Fields

9.1.1 Rings Defined

The simplest structure with two binary operations—and therefore the point where we begin—is called a *ring*. Because the assumptions we make about ring operations so closely resemble those for addition and multiplication on the integers, it is common to use the notations $+$ and either $\cdot$ or juxtaposition for the two operations, and we will shamelessly call them addition and multiplication, respectively, even though by doing so we run some slight risks. One risk is that we might inadvertently think that some of the rings we create are more like the integers with their addition and multiplication than the assumptions justify, because the WOP makes the integers a very special kind of ring. So we have to be careful not to bring any excess baggage from our understanding of the integers that the general ring assumptions do not imply. On the other hand, being able to envision the integers as a sort of quintessential example of a ring means that some of the results we proved for the integers will translate directly over to any ring and be clear to us right away. So some of the theorems we will state in this section will require very little in the way of a new proof but will mimic those for the integers with very little or no variation. The second risk we run in using notation already associated with the integers is that it might stifle our imagination when we try to create new and interesting rings. Creative minds have concocted quite a number of very interesting rings from interesting sets and definitions of equality, addition, and multiplication. We will see several of them. So here is the definition of ring, along with an enumeration of its defining characteristics.

Definition 9.1.1 Suppose R is a non-empty set endowed with binary operations $+$ (addition) and $\cdot$ (multiplication), such that R is an abelian group under addition, with additive identity 0 and additive inverse operation $-$, and multiplication is an associative binary operation that distributes over addition from the left and right. The algebraic structure $(R, +, 0, -, \cdot)$ is called a *ring*. If multiplication is a commutative operation, then R is called a *commutative ring*. If there exists a *nonzero* identity element for multiplication, we denote such an element e, and it is called a *unity element* or simply a *unity*.

Let's spell out the essential features of a ring in glorious detail.

(R1) Addition is well defined (G1).

(R2) Addition is closed (G2).

(R3) Addition is associative (G3).

(R4) There exists $0 \in R$ such that $a + 0 = a$ for all $a \in R$ (G4).

(R5) For all $a \in R$, there exists $-a \in R$ such that $a + (-a) = 0$ (G5).

(R6) Addition is commutative (abelian).

(R7) Multiplication is well defined (G1).

(R8) Multiplication is closed (G2).

(R9) Multiplication is associative (G3).

(R10) For all $a, b, c \in R$, $a(b + c) = ab + ac$ and $(b + c)a = ba + ca$.

Definition 9.1.1 requires that, if a ring has a unity element, it must be nonzero. Otherwise, the distributive property causes the entire ring to collapse to $\{0\}$, just as it does in the real numbers if $1 = 0$. Since the trivial ring $\{0\}$ is not particularly exciting in its internal structure, and since an occasional general result about rings with unity would not apply to the trivial ring, we simply insist for convenience that e is nonzero.

Definition 9.1.1 does not mention the existence of multiplicative inverses. Certainly, if a ring has a unity, some elements other than the unity might have a multiplicative inverse. We will address this issue in Section 9.3.

We can only hope that a lot of our examples of groups and their binary operations can serve as raw material from which to construct rings, or at least that new settings we will create can exploit binary operations whose basic features are transparent enough for us to make our way through verification of properties R1–R10 quick and relatively painless. Here are a few examples of rings. We will return to all of them several times in later sections.

Example 9.1.2 The integers with addition and multiplication form a commutative ring with unity element 1. Similarly, the rational numbers, real numbers, and complex numbers under the same operations are commutative rings with unity. The even integers are a commutative ring without unity. ■

EXERCISE 9.1.3 For $n \geq 2$, consider $\mathbb{Z}_n$ with addition $\oplus_n$ defined as in Eq. (8.39). Define multiplication $\otimes_n$ in an analogous way:

$$[(n) + a] \otimes_n [(n) + b] = (n) + ab \tag{9.1}$$

Table 8.2 illustrates the behavior of multiplication in $\mathbb{Z}_6$. From all our work in Section 8.3 on $\mathbb{Z}_n$ as an abelian additive group, properties R1–R6 are satisfied. Show the following for $\mathbb{Z}_n$.

(a) Addition and multiplication satisfy properties R7–R10.

(b) Multiplication is commutative.

(c) There exists a multiplicative identity.

As we will see later, the value of n has much to do with which elements of $\mathbb{Z}_n$ have a multiplicative inverse.

It would be a shame not to introduce an example of a matrix ring at this point, for matrices are very important in both theoretical and applied mathematics. They are the bread and butter of linear algebra and of many highly computational processes that have become practical only since the advent of computers. For our purposes, they are a storehouse of examples that exhibit all kinds of interesting behaviors.

Example 9.1.4 A *matrix* is a two-dimensional array, typically of real numbers:

$$A = \begin{bmatrix} a_{1,1} & a_{1,2} & \cdots & a_{1,n} \\ a_{2,1} & a_{2,2} & \cdots & a_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m,1} & a_{m,2} & \cdots & a_{m,n} \end{bmatrix} \tag{9.2}$$

The matrix in Eq. (9.2) is said to have dimensions m by n, and the set of all $m \times n$ matrices with real number entries is denoted $\mathbb{R}_{m\times n}$. If a matrix A has dimensions $m \times n$, we sometimes write it as $A_{m\times n}$ if we need to display its dimensions. Notice how entries of A are tagged with a row and column number and in that order. The commas in the subscripts are annoying to write and probably not necessary, so we will omit them if we can get away with it and still be clear. Thus a_{42} means the entry down in row 4 and across in column 2. Most of the matrices we will use will be 2×2, because we do not need to be any more general than that to construct some interesting examples. ■

Before we even discuss binary operations on sets of matrices, we need to define what it means for two matrices to be equal. We say two matrices, $A_{m_1\times n_1}$ and $B_{m_2\times n_2}$, are equal if two conditions are satisfied. First, the dimensions of A must be the same as those of B; that is, $m_1 = m_2$ and $n_1 = n_2$. Second, all their corresponding entries must equal as real numbers: $a_{jk} = b_{jk}$ for all j and k. Clearly, this definition is an equivalence relation, for the fact that properties E1–E3 are satisfied for the dimensions and all the real number entries in the matrices involved reveals that our definition of matrix equality also satisfies properties E1–E3.

The addition of matrices requires that they have the same dimensions, and $A + B$ is merely the matrix whose entries are the sums of the corresponding entries from A and B. For the 2×2 case, writing $A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}$ and $B = \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix}$, we have

$$A + B = \begin{bmatrix} a_{11} + b_{11} & a_{12} + b_{12} \\ a_{21} + b_{21} & a_{22} + b_{22} \end{bmatrix} \tag{9.3}$$

Matrix multiplication is more complicated than addition. In general, in order for the product of two matrices to be defined, they do not have to have the same dimensions, but some relationship between their dimensions must be satisfied. Since two square matrices $(n \times n)$ can always be multiplied to produce another $n \times n$ matrix, we will define multiplication only for this case. The 2×2 case should get the point across, though we will explain it in more general language. Writing

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} e & f \\ g & h \end{bmatrix} \tag{9.4}$$

we define AB in the following way. To calculate entry (j, k) (row j, column k) in the product, mentally highlight row j of A and column k of B. Mentally run your fingers across row j of A and down column k of B, multiplying the pairs $a_{j1}b_{1k}$, $a_{j2}b_{2k}$, etc; and add up these products. This sum of product pairs is entry (j, k) in AB. So for A and B in (9.4),

$$AB = \begin{bmatrix} ae + bg & af + bh \\ ce + dg & cf + dh \end{bmatrix} \tag{9.5}$$

Exercise 9.1.5 Let $A = \begin{bmatrix} 2 & -1 \\ 0 & 5 \end{bmatrix}$ and $B = \begin{bmatrix} 4 & 2 \\ 4 & 1 \end{bmatrix}$. Calculate AB and BA.

By Exercise 9.1.5, matrix multiplication is not commutative. With these definitions of matrix addition and multiplication, consider

$$\mathbb{R}_{2\times 2} = \left\{ \begin{bmatrix} a & b \\ c & d \end{bmatrix} : a, b, c, d \in \mathbb{R} \right\} \tag{9.6}$$

We show $\mathbb{R}_{2\times 2}$ is a noncommutative ring with unity element. That matrix addition is well defined (R1) is transparent and notationally tedious. If $A = B$ and $C = D$, then showing $A + C = B + D$ amounts to nothing more than applying the fact that real number addition is well defined to each entry in $A + C$ and $B + D$. Similar drudgery reveals that matrix multiplication is well defined, so property R7 holds. Closure of addition and multiplication in $\mathbb{R}_{2\times 2}$ is immediate from the definitions, so that properties R2 and R8 hold. Associativity and commutativity of addition in the real numbers mean that properties R3 and R6 hold. The matrix denoted $\mathbf{0}$, with all zero entries, is the additive identity, and

$$-\begin{bmatrix} a & b \\ c & d \end{bmatrix} = \begin{bmatrix} -a & -b \\ -c & -d \end{bmatrix} \tag{9.7}$$

The only remaining properties are R9 and R10, associativity of multiplication and left and right distributivity. These are not immediately obvious, but verifying them with all the necessary notation is about as exciting as counting stripes on the highway, and you are about as likely to make a mistake along the way.

So $\mathbb{R}_{2\times 2}$ is a noncommutative ring. Does it have a unity element? Why yes it does. Writing

$$I_{n\times n} = \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{bmatrix} \tag{9.8}$$

the matrix with 1s down the *main diagonal* and 0s elsewhere, you can see that $I_{2\times 2}A = AI_{2\times 2} = A$ for all $A \in \mathbb{R}_{2\times 2}$. Finally, by exactly the same reasoning we have used here, $\mathbb{R}_{n\times n}$ is a noncommutative ring with unity element $I_{n\times n}$, as are $\mathbb{Q}_{n\times n}$ and $\mathbb{Z}_{n\times n}$.

Here is an example of a ring that we can create from two given rings R and S. Verifying that this creation is a ring is both tedious and transparent, but you should at least mentally walk through the steps (or at least the first few of them) to see that it satisfies properties R1–R10. Rather than going crazy with notation to distinguish operations and elements of R from those of S, writing expressions like 0_R and $+_S$, we will assume your acquired level of mathematical sophistication makes them unnecessary. Just make sure you notice which set all the operations are being performed in.

Example 9.1.6 Suppose R and S are rings, and consider

$$R \times S = \{(r, s) : r \in R, s \in S\} \tag{9.9}$$

the Cartesian product of R and S. Define $(r_1, s_1) = (r_2, s_2)$ in $R \times S$ if $r_1 = r_2$ and $s_1 = s_2$. Define addition and multiplication in $R \times S$ by

$$(r_1, s_1) + (r_2, s_2) = (r_1 + r_2, s_1 + s_2) \quad \text{and} \quad (r_1, s_1) \cdot (r_2, s_2) = (r_1 r_2, s_1 s_2) \tag{9.10}$$

Then $R \times S$ is a ring under the operations defined in Eqs. (9.10). The zero element of $R \times S$ is $(0, 0)$, and $-(r, s) = (-r, -s)$. If R and S each have a unity element, then so does $R \times S$, and $e_{R\times S} = (e_R, e_S)$. ■

Exercise 9.1.7 Let R be the set of all functions defined on the real numbers. We use Definition 4.1.9 as our definition of equality. For two functions f and g, define their sum $f + g$ by the rule $(f + g)(x) = f(x) + g(x)$, and let composition serve as multiplication. Show that R with these operations is a non-commutative *near ring* with unity, in that one item in the list R1–R10 fails.

Exercise 9.1.8 Let U be a non-empty set and $\mathcal{P}(U)$ be the family of all subsets of U. Using symmetric difference as addition and intersection as multiplication, show that $\mathcal{P}(U)$ is a commutative ring with unity.

9.1.2 Fields Defined

It might seem strange to introduce our next term at this point, but it turns out to be more convenient as we progress through the theory of rings. A *field* is a special kind of ring, and its defining characteristics make it the most specialized kind of ring we will study in this text. Although we won't delve deeply into a general theory of fields, we will notice a few of their characteristics that are pretty easy to pick up along our way.

Definition 9.1.9 Suppose K is a commutative ring with unity, with the property that every nonzero element has a multiplicative inverse. Then K is called a *field*.

In addition to properties R1–R10, a field K must have the following features.

(K11) There exists a nonzero element e such that $e \cdot k = k$ for all $k \in K$ (G4).

(K12) For all nonzero $k \in K$, there exists k^{-1} such that $k \cdot k^{-1} = e$ (G5).

(K13) Multiplication is commutative (abelian).

Notice that properties K11–K13 complete the requirements for $K^\times$ to be an abelian group under multiplication. Thus a shorthand way of defining a field K is to say that K is an abelian group under addition and that $K^\times$ is an abelian group under multiplication.

Example 9.1.10 The rational numbers and the real numbers are fields. Also, the complex numbers are a field. In Section 8.1, we showed that $\mathbb{C}$ with addition is an abelian group, and in Exercise 8.1.20, you showed that $\mathbb{C}^\times$ with multiplication is an abelian group. ■

Example 9.1.11 $\mathbb{R}_{2\times2}$ is not a field because multiplication is not commutative. ■

Exercise 9.1.12 Not only is multiplication in $\mathbb{R}_{2\times2}$ not commutative, but also many elements do not have a multiplicative inverse. Show that $\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$ has no inverse by showing that the equation

$$\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}\begin{bmatrix} a & b \\ c & d \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \tag{9.11}$$

has no solution $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$.

9.2 Subrings

Suppose R is a ring and S is a subset of R. If S is also a ring under the same operations, we say that S is a *subring* of R. Demonstrating S is a subring, some of the properties R1–R10 are inherited from R, while some (the closure properties) must be shown for S.

(S1) S is closed under addition (R2, H1).

(S2) S contains the additive identity (R4, H2).

(S3) S is closed under additive inverses (R5, H3).

(S4) S is closed under multiplication (R8, H1).

A subring S of a ring R is merely a subgroup of the additive group that is also closed under multiplication. We call $\{0\}$ the *trivial* subring, and all subrings other than $\{0\}$ and R itself are called *proper* subrings.

Example 9.2.1 The integers are a subring of the rationals, and the rationals are a subring of the real numbers. ■

If R has a unity, it is not necessary that the unity be in S in order for S to be a subring of R.

Example 9.2.2 Let m be a positive integer, and write $m\mathbb{Z} = \{km : k \in \mathbb{Z}\}$, the set of integer multiples of m. This is another common notation for the set in Eq. (8.31), where we denoted by (m) the subgroup of the additive group of integers generated by m. Then $m\mathbb{Z}$ is a subring of the integers, because it is a subgroup of the integers under addition, and it is closed under multiplication. If $m \geq 2$, then $m\mathbb{Z}$ does not contain 1. ■

EXERCISE 9.2.3 Find all subrings of $\mathbb{Z}_6$ and $\mathbb{Z}_7$.

Example 9.2.4 $\mathbb{Z}_{2\times 2}$ is a subring of $\mathbb{R}_{2\times 2}$. ■

EXERCISE 9.2.5 Call a square matrix *diagonal* if its only nonzero entries lie on the main diagonal. Let $D_{2\times 2}$ be the subset of $\mathbb{Z}_{2\times 2}$ consisting of the diagonal matrices. Then $D_{2\times 2}$ is a subring of $\mathbb{Z}_{2\times 2}$.

Example 9.2.6 In this example, we create a subring of the rational numbers that we will return to in Section 9.9. Let $\mathbb{Q}_{OD}$ be the subset of the rational numbers whose denominators are odd. There is more than one way to denote an element of $\mathbb{Q}_{OD}$. The form

$$\left\{\frac{m}{2n+1} : m, n \in \mathbb{Z}\right\} \tag{9.12}$$

is an obvious way. But another useful way to denote the set is to exploit the prime factorization of the numerator and denominator, isolating 2 to keep it separate from all the other primes involved. This works for all elements except zero, which we throw in separately.

$$\begin{aligned}\mathbb{Q}_{OD} = \{0\} \cup \Bigg\{ \pm\frac{2^n p_1 p_2 \cdots p_r}{q_1 q_2 \cdots q_s} &: n \in \mathbb{W} \text{ and } p_i, q_i \text{ are } \textit{odd} \text{ primes} \\ &\text{for all } 1 \leq i \leq r \text{ and } 1 \leq i \leq s \Bigg\}\end{aligned} \tag{9.13}$$

The amount of repetition among the p_i and q_i does not matter. And notice that the form of an element in Eq. (9.13) includes 1 by letting $n = 0$, $r = s = 1$ and

$p_1 = q_1 = 3$. Verifying that $\mathbb{Q}_{OD}$ has properties S1–S4 is quick. Using either form (9.12) or (9.13), we see that both addition and multiplication are closed because the product of denominators of two elements is odd. That $\mathbb{Q}_{OD}$ contains 0 and additive inverses is obvious. Thus $\mathbb{Q}_{OD}$ is a subring of the rationals. ■

EXERCISE 9.2.7 Show that the set of all rational numbers with even numerator and odd denominator is a subring of $\mathbb{Q}_{OD}$.

Example 9.2.8 Let R be the near ring in Example 9.1.7, and let S be the set of all polynomial functions. Although it is notationally tedious to verify, S satisfies properties S1–S4 with respect to R. For the sum of two polynomials is a polynomial, the zero function $\mathbf{0}(x) = 0$ is a polynomial function whose coefficients are all zero, and the additive inverse of a polynomial is a polynomial. Furthermore, the composition of two polynomials is a polynomial. Note that S contains the unity element, which is the identity function. ■

EXERCISE 9.2.9 Let U be a non-empty set, and let A be a subset of U. Show that $\mathcal{P}(A)$ is a subring of $\mathcal{P}(U)$ from Exercise 9.1.8.

If you are required to show that a given set S is a subring of R, you might save yourself some time if you exploit the fact that S1–S3 are merely the subgroup properties H1–H3 applied to R and S as an abelian additive group and subgroup. If you have already shown that $(S, +, 0, -)$ is a subgroup of $(R, +, 0, -)$, then a lot of your work in showing S is a subring of R is already done. Keep that in mind as you prove the following.

EXERCISE 9.2.10 Suppose $\mathcal{F}$ is a family of subrings of a ring R. Then $\cap_{S\in\mathcal{F}} S$ is a subring of R.

Although subrings are very interesting in their own right, they are bland when compared to the special kind of subring called an *ideal*. Ideals come in all kinds of interesting flavors, and we will taste several in Section 9.5.

EXERCISE 9.2.11 Define e_R to be a right unity for a ring R if $r \cdot e_R = r$ for all $r \in R$. Similarly, call e_L a left unity if $e_L \cdot r = r$ for all $r \in R$. Let

$$L = \left\{ \begin{bmatrix} a & 0 \\ b & 0 \end{bmatrix} : a, b \in \mathbb{R} \right\} \tag{9.14}$$

(a) Show L is a ring by showing it is a subring of $\mathbb{R}_{2\times 2}$.

(b) Find an element of L that is a right unity but not a left unity.

(c) Show that this right unity is not unique.

9.3 Ring Properties

Since a ring is an abelian group under its addition operation, all properties of abelian groups you proved in Chapter 8 apply to addition. With regard to multiplication and its interaction with addition, it would probably be a good idea to swing back through Chapters 2 and 3 and point out theorems that we proved for the real numbers that exploited only their ring properties. A lot of theorems will then translate directly over to a general ring. The only difference is that multiplication might not be commutative, so we have to state and prove certain theorems in two-sided language to get the full strength. We will state the theorems here, with appropriate comments along the way. You will prove some of them as exercises. The corollaries should be mere observations.

Exercise 9.3.1 If R is a ring, then $a \cdot 0 = 0 \cdot a = 0$ for all $a \in R$.

Exercise 9.3.1 implies that zero will not have a multiplicative inverse in a ring with unity.

Exercise 9.3.2 If R is a ring and $a, b \in R$, then

(a) $(-a)b = -(ab)$

(b) $a(-b) = -(ab)$

(c) $(-a)(-b) = ab$

Corollary 9.3.3 If R is a ring with unity e, then $(-e)a = a(-e) = -a$ for all $a \in R$.

In an abelian group with operation $*$, $(a * b)^{-1} = a^{-1} * b^{-1}$. Since a ring is an abelian group under addition, this translates to the following additive form.

Theorem 9.3.4 If R is a ring, then $-(a + b) = (-a) + (-b)$ for all $a, b \in R$.

The distributive property extends nicely in a general ring to yield the following result analogous to Exercise 3.4.15 and Theorem 3.4.16.

Theorem 9.3.5 If R is a ring and $a, b_1, b_2, \ldots b_n \in R$, then

$$a \sum_{k=1}^{n} b_k = \sum_{k=1}^{n} (ab_k) \tag{9.15}$$

Theorem 9.3.6 If R is a ring and $a_1, a_2, \ldots, a_m, b_1, b_2, \ldots b_n \in R$, then

$$\left(\sum_{j=1}^{m} a_j\right)\left(\sum_{k=1}^{n} b_k\right) = \sum_{j=1}^{m}\left(\sum_{k=1}^{n} a_j b_k\right) \tag{9.16}$$

As in a group, if a ring has a unity element, it can have only one. Your proof from Exercise 8.1.23 would translate directly over to a general ring, pretty much word for word.

Theorem 9.3.7 If R is a ring with unity, then the unity element is unique.

Even though elements of a ring with unity are not assumed to have multiplicative inverses, some of them might. If for a given x there exists y such that $xy = yx = e$, then x is called a *unit* of the ring. Notice that Exercise 9.3.1 implies that zero is not a unit.

Exercise 9.3.8 Find, with verification, all units in the following rings.

(a) $\mathbb{Z}$

(b) $\mathbb{Z}_{12}$

(c) $\mathbb{Z}_7$

(d) $D_{2\times 2}$ from Exercise 9.2.5

(e) The near ring of functions from Exercise 9.1.7

(f) The power set ring from Exercise 9.1.8

Divisibility in a ring is defined in pretty much the same way it is in the integers, except that we must distinguish between left and right divisors.

Definition 9.3.9 Let R be a ring, and let a and b be elements of R, where a is nonzero. Then a is called a *left divisor* of b provided there exists *nonzero* $k \in R$ such that $ak = b$. Similarly, a is called a *right divisor* of b if there exists nonzero $k \in R$ such that $ka = b$. If R is a commutative ring and there exists nonzero $k \in R$ such that $ak = b$, we say simply that a is a *divisor* of b, or that *a divides b*, and we write $a \mid b$.

If a divides b, it does not necessarily mean that the k such that $ak = b$ or $ka = b$ is unique. In a more specialized ring we will study in Section 9.8, however, we will have uniqueness.

Exercise 9.3.10 Find *nonzero* elements a, b, k_1, k_2 in each of the following rings where $ak_1 = b$ and $ak_2 = b$, but k_1 is different from k_2.

(a) $\mathbb{Z}_{12}$

(b) $\mathbb{Z}_{2\times 2}$

Having defined divisors, we can now define what it means for an element of a ring to be prime, though we will not really look into any of its properties until Section 9.8. To motivate the term by returning to the positive integers, a prime p has exactly two distinct positive integer divisors, 1 and p. Thus by definition, 1 is not prime. So if p is a prime number and $p = ab$ is any factorization of p into positive integers a and b, then either $a = 1$ or $b = 1$. In the ring of integers, primes are extended to include the negatives of the prime natural numbers. Thus in the integers, for p to be prime means that if $p = ab$ is any factorization of p into integers a and b, then either a or b is ± 1. In Exercise 9.3.8(a), you showed that the integer units are ± 1. In a general ring, a prime element is defined by using the language of divisors and units.

Definition 9.3.11 Suppose R is a ring, and p is a ring element that is not a unit. Then p is said to be prime provided every factorization $p = ab$ into ring elements a and b implies either a or b is a unit.

Definition 9.3.11 automatically excludes zero from being prime, for $0 = 0 \cdot 0$, and zero is not a unit.

Exercise 9.3.12 Find all primes in $\mathbb{Z}_4, \mathbb{Z}_6$, and $\mathbb{Z}_7$.

In Exercises 2.1.8 and 2.1.19, you proved multiplicative cancellation for nonzero real numbers and the principle of zero products. In a ring, these properties do not necessarily apply, either as part of the definition or as logical consequences of it. However, these properties will reappear in Section 9.8 when we discuss integral domains. From Definition 9.3.9, if $ab = 0$ while neither a nor b is zero, then a and b are called *divisors of zero* or *zero divisors*.

Exercise 9.3.13 Find a zero divisor in each of the following rings.

(a) $\mathbb{Z}_n$ for some strategically chosen n

(b) $\mathbb{Z}_{2\times 2}$

Exercise 9.3.14 Let R and S be rings. Find all units and zero divisors in $R \times S$, in terms of the units and zero divisors of R and S individually.

Exercise 9.3.15 Find all zero divisors in the power set ring from Exercise 9.1.8.

Exercise 9.3.16 Show that if a is a divisor of zero in a ring with unity, then a is not a unit.

Exercise 9.3.17 If R is a ring with unity and x is a unit, then the element y satisfying $xy = yx = e$ is unique.

A lot of the results in Section 2.2 involved ordering of real numbers as measured by $<$. Since a ring does not necessarily have any such way of comparing its elements, none of these results has meaning in a ring without such a basis for comparison being defined. Thus rings do not necessarily have positive and negative elements; there is not necessarily a way to measure the size of elements as absolute value does in the real numbers, and there is not necessarily a way to make the WOP applicable to subsets. And just because the equation $x^2 = -1$ has no real solution, it does not mean that the equation $x^2 = -e$ will not have a solution.

Exercise 9.3.18 Determine all solutions to $x^2 = -1$ in $\mathbb{Z}_5, \mathbb{Z}_6, \mathbb{Z}_7$, and $\mathbb{Z}_{10}$.[1]

Since a ring with its addition operation is an abelian group, we can apply the additive forms of the equations from Exercise 8.2.16.

$$0a = 0 \tag{9.17}$$

$$(n+1)a = na + a \tag{9.18}$$

$$(-n)a = -(na) \tag{9.19}$$

$$ma + na = (m+n)a \tag{9.20}$$

$$m(na) = (mn)a \tag{9.21}$$

$$n(a+b) = na + nb \tag{9.22}$$

Be careful to distinguish between integers and ring elements in these equations. Also, note that Eqs. (9.20)–(9.22) are not the associative and distributive properties that characterize either the ring or the integers individually.

Example 9.3.19 In $\mathbb{Z}_{2\times 2}$, let $A = \begin{bmatrix} 1 & 2 \\ -1 & 0 \end{bmatrix}$ and $B = \begin{bmatrix} 5 & 0 \\ 2 & 2 \end{bmatrix}$. Then

$$0A = 0\begin{bmatrix} 1 & 2 \\ -1 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} \tag{9.23}$$

$$6A = \overbrace{\begin{bmatrix} 1 & 2 \\ -1 & 0 \end{bmatrix} + \cdots + \begin{bmatrix} 1 & 2 \\ -1 & 0 \end{bmatrix}}^{6\text{ times}} = \begin{bmatrix} 6 & 12 \\ -6 & 0 \end{bmatrix} \tag{9.24}$$

$$-6A = -(6A) = -\begin{bmatrix} 6 & 12 \\ -6 & 0 \end{bmatrix} = \begin{bmatrix} -6 & -12 \\ 6 & 0 \end{bmatrix} \tag{9.25}$$

$$\begin{aligned} 4A + 2A &= 4\begin{bmatrix} 1 & 2 \\ -1 & 0 \end{bmatrix} + 2\begin{bmatrix} 1 & 2 \\ -1 & 0 \end{bmatrix} = \begin{bmatrix} 4 & 8 \\ -4 & 0 \end{bmatrix} + \begin{bmatrix} 2 & 4 \\ -2 & 0 \end{bmatrix} = \begin{bmatrix} 6 & 12 \\ -6 & 0 \end{bmatrix} \\ &= 6\begin{bmatrix} 1 & 2 \\ -1 & 0 \end{bmatrix} = 6A = (4+2)A \end{aligned} \tag{9.26}$$

$$4(2A) = 4\begin{bmatrix} 2 & 4 \\ -2 & 0 \end{bmatrix} = \begin{bmatrix} 8 & 16 \\ -8 & 0 \end{bmatrix} = 8\begin{bmatrix} 1 & 2 \\ -1 & 0 \end{bmatrix} = 8A = (4 \cdot 2)A \tag{9.27}$$

$$\begin{aligned} 4(A+B) &= 4\left(\begin{bmatrix} 1 & 2 \\ -1 & 0 \end{bmatrix} + \begin{bmatrix} 5 & 0 \\ 2 & 2 \end{bmatrix}\right) = 4\begin{bmatrix} 6 & 2 \\ 1 & 2 \end{bmatrix} = \begin{bmatrix} 24 & 8 \\ 4 & 8 \end{bmatrix} = \begin{bmatrix} 4 & 8 \\ -4 & 0 \end{bmatrix} + \begin{bmatrix} 20 & 0 \\ 8 & 8 \end{bmatrix} \\ &= 4\begin{bmatrix} 1 & 2 \\ -1 & 0 \end{bmatrix} + 4\begin{bmatrix} 5 & 0 \\ 2 & 2 \end{bmatrix} = 4A + 4B \end{aligned} \tag{9.28}$$

■

[1] Calculate all values of x^2 and see if -1 is ever produced.

Exercise 9.3.20 Suppose R is a ring with unity element e. Show that $(me)(ne) = (mn)e$ for all positive integers m and n.

Even though multiplication in a ring is not necessarily accompanied by an identity and inverses for elements, we can use the multiplicative forms of the definitions of a^n in a limited way. For a ring element a, we begin by defining $a^1 = a$ and $a^{n+1} = a^n \cdot a$ for $n \geq 1$. If R has a unity element e, we also define $a^0 = e$, but only for nonzero a. If a is a unit of R, we can define $a^{-n} = (a^{-1})^n$.

Example 9.3.21 In $\mathbb{Z}_{2\times 2}$, let $A = \begin{bmatrix} 1 & 2 \\ -1 & 0 \end{bmatrix}$ Then

$$A^0 = I = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

$$A^1 = A = \begin{bmatrix} 1 & 2 \\ -1 & 0 \end{bmatrix}$$

$$A^2 = A \times A = \begin{bmatrix} -1 & 2 \\ -1 & -2 \end{bmatrix}$$

$$A^3 = A^2 \times A = \begin{bmatrix} -3 & -2 \\ 1 & -2 \end{bmatrix}, \text{ etc.}$$ ■

Exercise 9.3.22 Evaluate the first few (multiplicative) powers of 3 and 5 in $\mathbb{Z}_{10}$.

With these definitions of a^n for appropriate integers n, and by arguments exactly like those in Exercise 3.5.4, we have the following.

Theorem 9.3.23 Suppose R is a ring, $a, b \in R$, and let m and n be positive integers. Then

$$a^m \cdot a^n = a^{m+n} \tag{9.29}$$

$$(a^m)^n = a^{mn} \tag{9.30}$$

If R has a unity element and a is a nonzero ring element, Eqs. (9.29) and (9.30) hold for all nonnegative integers m and n. If a is a unit, these equations hold for all integers m and n. Furthermore, if R is commutative,

$$(ab)^n = a^n b^n \tag{9.31}$$

for all n for which a^n and b^n are defined.

One big difference in the way we visualize the integers and $\mathbb{Z}_n$ is that the integers extend out indefinitely along the number line in both directions, whereas $\mathbb{Z}_n$ is circular. If we generate both of these rings by considering $1, 1+1, 1+1+1$, and so on, no expression of the form $n1 = \sum_{k=1}^{n} 1$ ever produces a sum of zero in the integers, but $n1 = \sum_{k=1}^{n} 1 = 0$ in $\mathbb{Z}_n$. In a general ring with unity element e, whether or not some expression $ne = \sum_{k=1}^{n} e = 0$ ever occurs motivates a term.

Definition 9.3.24 Suppose R is a ring with unity, and suppose there exists a positive integer n such that $ne = 0$. Then the smallest such n for which this holds is called the *characteristic* of R, and is denoted char R. If no such positive integer exists, then R is said to have characteristic zero.

In $\mathbb{Z}_n$ the fact that $n \equiv_n 0$, or $n1$ is zero in $\mathbb{Z}_n$ means that char $\mathbb{Z}_n \leq n$. On the other hand, if m is a positive integer and $m \equiv_n 0$, that is, if $m1$ is zero in $\mathbb{Z}_n$, then m is a multiple of k, so that $m \geq n$. Thus char $\mathbb{Z}_n \geq n$, so that we have proved the following theorem.

Theorem 9.3.25 If $n \geq 2$, then char $\mathbb{Z}_n = n$.

Exercise 9.3.26 What is char$(\mathbb{Z}_4 \times \mathbb{Z}_{18})$? Explain.

If adding the unity element to itself n times produces a sum of zero, then the same is true for all elements of the ring.

Exercise 9.3.27 Let R be a ring with unity and nonzero characteristic n. Then $nx = 0$ for all $x \in R$.

9.4 Ring Extensions

We can create a very important type of algebraic structure from a given algebraic structure by tossing in a new element, stirring well, and letting the mixture expand into another algebraic structure of the same type. It is called the process of *adjoining* an element in order to create what is called an *extension* of the original structure. In this section we want to get acquainted with the creation of extensions by adjoining elements to *commutative* rings. In principle, only two types of ring extensions can result from adjoining an element. We will begin with a very specific example of the first type, but instead of building it up as an extension of a certain ring in the most rigorous way, we will just lay the whole structure out there, define equality and the operations, and show that what we have presented is a ring. But don't worry. We will make up for our lax introduction of this ring in Section 9.14, where we will see a more rigorous way to construct it. After we have presented our example of the first type of ring extension, we will point out how other extensions of the same type can be created through precisely the same reasoning. Finally, we construct the canonical example of the second type. We will use these constructions over and over throughout the rest of this chapter.

9.4.1 Adjoining Roots of Ring Elements

Example 9.4.1 Let $S = \{a + b\sqrt[3]{2} + c\sqrt[3]{4} : a, b, c \in \mathbb{Z}\}$, the set of all integer linear combinations of $\{1, \sqrt[3]{2}, \sqrt[3]{4}\}$. First, we define $x = a_1 + b_1\sqrt[3]{2} + c_1\sqrt[3]{4}$ to be equal to $y = a_2 + b_2\sqrt[3]{2} + c_2\sqrt[3]{4}$ provided $a_1 = a_2$, $b_1 = b_2$, and $c_1 = c_2$. Notice that this definition is an equivalence relation because it is just an application of integer equality in triplicate.[2] Define addition $\oplus$ and multiplication $\odot$ in a natural way, based on the extended distributive property and the behavior of $\sqrt[3]{2}$ in the real numbers:

$$
\begin{aligned}
&(a + b\sqrt[3]{2} + c\sqrt[3]{4}) \oplus (d + e\sqrt[3]{2} + f\sqrt[3]{4}) = (a + d) + (b + e)\sqrt[3]{2} + (c + f)\sqrt[3]{4} \\
&(a + b\sqrt[3]{2} + c\sqrt[3]{4}) \odot (d + e\sqrt[3]{2} + f\sqrt[3]{4}) = \\
&\qquad (ad + 2bf + 2ce) + (ae + bd + 2cf)\sqrt[3]{2} + (af + be + cd)\sqrt[3]{4}
\end{aligned} \tag{9.32}
$$

Because $\oplus$ and $\odot$ in S have the familiar behavior we expect when viewed within the context of the real numbers, we could simply use $+$ and $\cdot$. Just remember that a single entity in S is of the form $a + b\sqrt[3]{2} + c\sqrt[3]{4}$ and includes partial forms like 1, $6\sqrt[3]{2}$, or $4 - \sqrt[3]{4}$ by letting certain coefficients be zero.

From Eqs. (9.32) and closure of integer addition and multiplication, the closure of $\oplus$ and $\odot$ is immediately obvious. Also, S contains the additive identity $0 + 0\sqrt[3]{2} + 0\sqrt[3]{4}$, additive inverses $(-a) + (-b)\sqrt[3]{2} + (-c)\sqrt[3]{4}$, and unity element $1 + 0\sqrt[3]{2} + 0\sqrt[3]{4}$. Furthermore, all remaining ring properties are assumed for the real numbers and are therefore inherited by S. Since $\odot$ is commutative (next exercise), S is a commutative ring with unity element. ■

EXERCISE 9.4.2 Show that $\odot$ in Example 9.4.1 is commutative.

The notation we use for the ring in Example 9.4.1 is $\mathbb{Z}[\sqrt[3]{2}]$, which is meant to denote that the ring of integers has had an additional element $\sqrt[3]{2}$ thrown in, or *adjoined*, as we say. Adjoining an element to an algebraic structure is obviously different from unioning it onto the set. Instead of merely tossing it in as one more additional element, we toss it in, then combine it by addition and multiplication with itself and all other elements to expand into a ring. Thus you can see that the presence of $\sqrt[3]{4}$ in $\mathbb{Z}[\sqrt[3]{2}]$ is necessary so that we have closure of multiplication: $(\sqrt[3]{2})^2 = \sqrt[3]{4}$. But the fact that $(\sqrt[3]{2})^3$ is an integer means that expressions of the form $a + b\sqrt[3]{2} + c\sqrt[3]{4}$ are all that are necessary. By similar reasoning to that in

[2] This definition of equality will raise all kinds of concerns in the mind of your professor because of something you will probably just assume without any basis. You probably think that our definition of equality here exactly coincides with equality in the real numbers, so that two expressions in S are equal if and only if they are equal as real numbers. Clearly, if $x = y$ as we have defined equality for S, then $x = y$ as real numbers. But just because $a_1 + b_1\sqrt[3]{2} + c_1\sqrt[3]{4} = a_2 + b_2\sqrt[3]{2} + c_2\sqrt[3]{4}$ as real numbers, we cannot conclude immediately that $a_1 = a_2$, $b_1 = b_2$, and $c_1 = c_2$. If you crank out $5{,}096{,}516{,}652 - 5184\sqrt[3]{2} + 91{,}047{,}715{,}794\sqrt[3]{4}$ and $2{,}669{,}624{,}714 + 130{,}936{,}500{,}093\sqrt[3]{2} - 11{,}347{,}811{,}196\sqrt[3]{4}$ on a TI-85 calculator, it appears they might be equal in the real numbers. They are not, and it is indeed true that equality in $\mathbb{R}$ implies equality in S. That is, if there is a way to write a real number in the form $a + b\sqrt[3]{2} + c\sqrt[3]{4}$, then there is only one way to do it. To prove this, you need to know more about $\sqrt[3]{2}$ and $\sqrt[3]{4}$ as real numbers and their relationship to each other in the context of the integers. The term is *linear independence*, and you will see it in linear algebra.

Example 9.4.1, we could begin with the integers (or the rationals), choose a positive integer n and an integer x, define equality, addition, and multiplication, and show all ring properties (plus commutativity) for

$$\mathbb{Z}[\sqrt[n]{x}\,] = \{a_0 + a_1\sqrt[n]{x} + a_2\sqrt[n]{x^2} + \cdots + a_{n-1}\sqrt[n]{x^{n-1}} : a_i \in \mathbb{Z}\} \tag{9.33}$$

The form of elements of $\mathbb{Z}[\sqrt[n]{x}]$ in Eq. (9.33) assumes that n is the smallest positive integer such that x^n is an integer, so that none of the terms $\sqrt[n]{x}, \ldots, \sqrt[n]{x^{n-1}}$ is an integer.

As another example, $\mathbb{Q}[\sqrt{5}] = \{a + b\sqrt{5} : a, b \in \mathbb{Q}\}$ fits the form of Eq. (9.33). There is no reason x in Eq. (9.33) cannot be negative, and $\mathbb{Z}[\sqrt{-1}] = \mathbb{Z}[i] = \{a + bi : a, b \in \mathbb{Z}\}$ is called the *Gaussian integers*. $\mathbb{Z}[\sqrt{-5}]$ turns out to be an interesting ring that we will look at in Section 9.8. Finally, $\mathbb{R}[i] = \mathbb{C}$. Notice that $\mathbb{Z}$ is the subring of $\mathbb{Z}[\sqrt[3]{2}]$ consisting of all elements of the form $a + b\sqrt[3]{2} + c\sqrt[3]{4}$ where $b = c = 0$.

For the class of rings in the next theorem, we need to have a handle on the units when we get to Section 9.8. It is important to work through the proof of the next theorem because you will need to mimic the algebraic manipulation in an exercise that follows.

Theorem 9.4.3 Suppose p is a positive prime integer. Then the only units in $\mathbb{Z}[\sqrt{-p}]$ are ±1.

Proof. Suppose $a + b\sqrt{-p}$ is a unit. We show that $b = 0$ and $a = \pm 1$ by supposing

$$(a + b\sqrt{-p})(c + d\sqrt{-p}) = 1 \tag{9.34}$$

and drawing conclusions about a, b, c, d. Multiplying out the terms in Eq. (9.34) and using the definition of equality in $\mathbb{Z}[\sqrt{-p}]$ yields

$$ac - pbd = 1 \quad \text{and} \quad ad + bc = 0 \tag{9.35}$$

Squaring each side of Eqs. (9.35) yields

$$a^2c^2 - 2pabcd + p^2b^2d^2 = 1 \quad \text{and} \quad a^2d^2 + 2abcd + b^2c^2 = 0 \tag{9.36}$$

Multiplying the second equation in (9.36) by p and adding the two equations yields

$$\begin{aligned} a^2c^2 + pa^2d^2 + p^2b^2d^2 + pb^2c^2 &= 1 \\ a^2(c^2 + pd^2) + pb^2(pd^2 + c^2) &= 1 \\ (a^2 + pb^2)(c^2 + pd^2) &= 1 \end{aligned} \tag{9.37}$$

Each factor in Eq. (9.37) is a positive integer because the components are squared and p is positive. Thus by Exercise 2.2.7(f),

$$a^2 + pb^2 = 1 \quad \text{and} \quad c^2 + pd^2 = 1 \tag{9.38}$$

Furthermore, since $p \geq 2$, it must be that $b = d = 0$, so that $a = \pm 1$. □

EXERCISE 9.4.4 Using i, j, k as in the quaternion group (Example 8.1.18), construct the ring extension $\mathbb{Z}[i, j, k]$, defining equality, addition, and multiplication, then showing that all ring properties R1–R10 are satisfied. Is it commutative?

EXERCISE 9.4.5 Show that $S = \{a + 0i + cj + 0k : a, c \in \mathbb{R}\}$ is a subring of the ring from Exercise 9.4.4.

EXERCISE 9.4.6 Finding the units in a ring amounts to solving the equation $xy = 1$. In $\mathbb{Z}_6$, the only units are 1 and 5, so the equation $xy \equiv_6 1$ implies $x, y \in \{1, 5\}$. Use this fact and the technique in the proof of Theorem 9.4.3 to find all 16 units in $\mathbb{Z}_6[\sqrt{2}]$.

EXERCISE 9.4.7 Find, with verification, all units in $\mathbb{Z}[i]$.

EXERCISE 9.4.8 Prove that the following commutative rings with unity are fields by showing that every nonzero element is a unit.

(a) $\mathbb{Q}[\sqrt{2}]$

(b) $\mathbb{Q}[i]$

9.4.2 Polynomial Rings

The other type of extension we want to create might seem fundamentally different from the previous ones, but the principle is the same. Its one notable difference is that the new element we adjoin is, in a sense, more foreign to the original ring than numbers like $\sqrt{-5}$ are to the integers. The relationship of $\sqrt{-5}$ to the integers is characterized by the fact that $(\sqrt{-5})^2 = -5$, which is an integer, or if you prefer, $(\sqrt{-5})^2 + 5 = 0$. Similarly, if $x = \sqrt{1 + \sqrt[3]{2}}$, then $(x^2 - 1)^3 - 2 = 0$. Thus, as with $\sqrt{-5}$, there is some way to manipulate x using only the ring elements and operations to produce zero. The term that describes this relationship of $\sqrt{-5}$ and $\sqrt{1 + \sqrt[3]{2}}$ to the integers is *algebraic*, and numbers that are not algebraic are called *transcendental*. For example, π is transcendental over the integers because there is no way to combine π and any finite set of integers using the ring operations a finite number of times to produce zero. Strict definitions of these terms will come in your later work in algebra. For now, we simply construct an example in which the symbol we adjoin is transcendental over the integers because we define the ring and the behavior of the symbol to make it so.

Let R be a commutative ring, and write $R[t]$ to mean the set of all polynomials in the variable t, where the coefficients are elements of R. That is,

$$R[t] = \{a_n t^n + a_{n-1} t^{n-1} + \cdots + a_1 t + a_0 : n \in \mathbb{W}, a_k \in R \text{ for all } k, \\ \text{and } a_n \neq 0 \text{ if } n \neq 0\} \tag{9.39}$$

We insist that a_n be nonzero for $n \neq 0$ because it would be silly to suggest that the highest power term of an element of $R[t]$ is $a_n t^n$ when its coefficient a_n wipes it out. However, if $n = 0$, we do want to allow for $a_0 = 0$, the zero polynomial. We address an arbitrary element of $R[t]$ as f, as if it were a function. We will not, however, write it as $f(t)$. The reason is that we are not interested in an element of $R[t]$ primarily as an expression into which we substitute numeric values for t, but more as a string of symbols whose behavior in the ring we are constructing involves the symbol t merely as a way to describe how elements of $R[t]$ add and multiply.[3] First, we define $f = a_m t^m + \cdots + a_0$ and $g = b_n t^n + \cdots b_0$ to be equal if $m = n$ and $a_k = b_k$ for all $0 \leq k \leq n$, which is clearly an equivalence relation. Concerning the binary operations on $R[t]$, define addition and multiplication of two elements in the familiar way of adding and multiplying two polynomials. Notationally, it is very ugly to state the definitions of addition and multiplication formally, but you're certainly familiar with the way they are done. Assuming this familiarity, let's check that all ring properties R1–R10 are satisfied.

First of all, the fact that R has properties R1–R3 and R6 means that $R[t]$ does, too. Letting $n = 0$ and $a_0 = 0$ reveals that $f = 0$ (viewed as a polynomial in $R[t]$ and not as a mere element of R) is the additive identity (R4). The existence of additive inverses in R makes property R5 clear. Considering the way polynomials multiply, properties R7–R10 call on all the similar properties of both addition and commutative multiplication in R. Showing these is a mess but not difficult. Furthermore, multiplication in $R[t]$ is commutative because R is commutative, and if R has a unity element e, the polynomial e is the unity element in $R[t]$. Thus $R[t]$ is a commutative ring and has a unity element if R does. The new ring $R[t]$ is called the *polynomial ring* over R.

If we were to write an element of $\mathbb{Z}[\sqrt[3]{2}]$ as $c(\sqrt[3]{2})^2 + b\sqrt[3]{2} + a$, we could say that elements of $\mathbb{Z}[\sqrt[3]{2}]$ are three-term polynomials in the symbol $\sqrt[3]{2}$. The fundamental difference between $\mathbb{Z}[\sqrt[3]{2}]$ and $R[t]$ is that the three-term polynomials $c(\sqrt[3]{2})^2 + b\sqrt[3]{2} + a$ are all that is necessary to have closure of the ring operations when $\sqrt[3]{2}$ is adjoined to the integers. The behavior of $\sqrt[3]{2}$—that is, the fact that $(\sqrt[3]{2})^3$ is an integer—means it is not necessary to have terms of the form $(\sqrt[3]{2})^n$ for $n \geq 3$. But in $R[t]$, polynomials can be of any *degree*, whatever that means.

9.4.3 Degree of a Polynomial

If $f \in R[t]$ is written as $f = a_n t^n + \cdots + a_0$, where a_n is nonzero, we define the *degree* of f to be n, and we write $\deg f = n$. This definition does not assign a degree to the zero polynomial, so we will not assign a degree to it. Some authors define $\deg 0 = -\infty$. Even though $-\infty$ is not a real number, this degree assignment can make theorems involving $\deg f$ hold for the zero polynomial. Instead of assigning a degree to the zero polynomial and making it a special case in theorems, we will understand that polynomials whose degree we are working with are always nonzero polynomials. In $R[t]$, the degree of a polynomial is a measure of its size,

[3] This is a lie. But for now, it's true. See Section 9.12.

like absolute value in the real numbers, even though its properties do not really jibe with the norm properties N1–N3 (page 54). But as we will see later, it gives us at least some way to apply the WOP to nonzero elements of $R[t]$.

Exercise 9.4.9 Suppose R is a commutative ring, and let f and g be nonzero polynomials in $R[t]$. Then the following holds.

(a) If $\deg f = \deg g = n$, then $\deg(f+g) \leq n$.

(b) If $\deg f > \deg g$, then $\deg(f+g) = \deg f$.

(c) $\deg(fg) \leq \deg f + \deg g$.

It is easy to see why part (a) of Exercise 9.4.9 is an inequality, for if $f = 2t^2 + 1$ and $g = -2t^2 + 4t$ are elements of $\mathbb{Z}[t]$, then $f + g = 4t + 1$. But the fact that part (c) is an inequality instead of an equation might seem strange. The existence of zero divisors in R allows for this funny behavior in $R[t]$. For example, in $\mathbb{Z}_6[t]$,

$$(2t^2+3)(3t+3) = 6t^3 + 6t^2 + 9t + 9 = 3t + 3 \tag{9.40}$$

So the degree of a product can be strictly less than the sum of the degrees of the factors. Equation (9.40) also illustrates that it is possible that one polynomial may divide another polynomial of lower degree. These unfamiliar idiosyncracies can happen because of the presence of zero divisors in R. In Section 9.8, these behaviors will go away when we look at the polynomial ring over an integral domain.

Exercise 9.4.10 Calculate $(2t^2 + 4t + 1)(3t^3 + 3t + 4)$ in $\mathbb{Z}_5[t]$, then in $\mathbb{Z}_6[t]$.

Exercise 9.4.11 Give an example of a ring R and a polynomial in $R[t]$ that is a divisor of zero.

9.5 Ideals

In the same way a normal subgroup is a special kind of subgroup that exhibits a characteristic stronger than closure, we define a special class of subring where we have something stronger than closure of multiplication.

Definition 9.5.1 Suppose R is a ring and I is a subring of R with the property that $rx \in I$ for all $x \in I$ and $r \in R$. Then I is called a *left ideal* of R. Similarly, if $xr \in I$ for all $x \in I$ and $r \in R$, then I is called a *right ideal*. If I is both a left and right ideal, it is called a *two-sided ideal*. If R is commutative, there is no distinction between left and right ideals, and we will simply use the term *ideal*. Also, for simplicity, if R is not commutative, then we will refer simply to an *ideal* to denote an arbitrary left or right ideal.

An ideal, say a left ideal, has properties S1–S4, but property S4 is replaced with the stronger property that $rx \in I$ for all $x \in I$ and $r \in R$. What makes a left ideal more than a subring is that it "absorbs" multiplication from the left, and similarly for a right ideal. Thus an ideal is similar to a normal subgroup in that its elements exhibit a certain behavior in the presence of *all* ring elements. Let's list the defining properties of a left (or right) ideal.

(Y1) I is closed under addition (S1).

(Y2) I contains the additive identity (S2).

(Y3) I is closed under additive inverses (S3).

(Y4) For all $x \in I$ and $r \in R$, $rx \in I$ (or $xr \in I$).

The ideal $\{0\}$ is called the *trivial* ideal. Also, a ring is an ideal of itself. All other ideals besides these are called *proper* ideals.

Example 9.5.2 In $\mathbb{Z}_{12}$, the set $\{0, 4, 8\}$ is an ideal, for sums and additive inverses of multiples of 4 are also multiples of 4, and any ring element times a multiple of 4 is a multiple of 4. ■

EXERCISE 9.5.3 Let S be the set of all polynomials in $\mathbb{Z}[t]$ whose constant term is even. Then S is an ideal in $\mathbb{Z}[t]$.

Example 9.5.4 In Example 9.2.2 we showed that $m\mathbb{Z}$ is a subring of the integers. If n is any integer and mk is an element of $m\mathbb{Z}$, then $(mk)n = m(kn)$, which is also an element of $m\mathbb{Z}$. Thus $m\mathbb{Z}$ is a right ideal. Since integer multiplication is commutative, $m\mathbb{Z}$ is simply an ideal. ■

Example 9.5.4 is a special case of a type of ideal in a ring with unity. Starting with any integer m, we multiply m by every integer (in this case on the right) to create an ideal. Similarly in an arbitrary ring, we may take any ring element and create a left or right ideal by multiplying it on the left or right by every element of the ring.

EXERCISE 9.5.5 Suppose R is a ring and $a \in R$. Then the set

$$Ra = \{ra : r \in R\} \tag{9.41}$$

is a left ideal in R. (By similar reasoning, $aR = \{ar : r \in R\}$ is a right ideal in R.)

EXERCISE 9.5.6 Construct $3\mathbb{Z}_7$ and $3\mathbb{Z}_{12}$.

If R does not have a unity element, then the ideals Ra and aR might not contain a.

Example 9.5.7 Let E be the ring of even integers. Then

$$6E = \{\ldots, -24, -12, 0, 12, 24, \ldots\} \tag{9.42}$$

is an ideal of E that does not contain 6. ■

Example 9.5.8 For the polynomial ring $\mathbb{Z}[t]$, $(3t+2)\mathbb{Z}[t]$ is the ideal of all multiples of $3t+2$ in $\mathbb{Z}[t]$. Since $\mathbb{Z}[t]$ has unity, $3t+2 \in (3t+2)\mathbb{Z}[t]$. ■

All the previous examples of ideals are in commutative rings. The next exercise illustrates an interesting possibility in a noncommutative ring.

EXERCISE 9.5.9 Let $M = \left\{ \begin{bmatrix} a & b \\ 0 & 0 \end{bmatrix} : a, b \in \mathbb{Z} \right\}$. Show that M is a right ideal in $\mathbb{Z}_{2\times 2}$ but not a left ideal.

A result similar to Exercise 9.2.10 holds for ideals, but we must distinguish between left and right ideals.

Theorem 9.5.10 Suppose $\mathcal{F}$ is a family of left (or right) ideals of a ring R. Then $\cap_{I\in\mathcal{F}} I$ is a left (or right) ideal of R.

EXERCISE 9.5.11 Prove the left-sided case of Theorem 9.5.10.

In a commutative ring, where there is no difference between left and right ideals, Exercise 9.5.11 says that the intersection of a family of ideals is an ideal. In the next section, you will show that the intersection of a left ideal and a right ideal need not be either a left or right ideal.

EXERCISE 9.5.12 Demonstrate a ring R and two ideals I_1 and I_2 such that $I_1 \cup I_2$ is not an ideal in R.

Although the union of two ideals is not necessarily an ideal, there is a theorem that will come in handy in Section 9.9 that says something about the union across a special family of ideals.

Theorem 9.5.13 Suppose $\{I_n\}_{n=1}^{\infty}$ is a family of left (or right) ideals of a ring with the property that $I_n \subseteq I_{n+1}$ for all n. Then $\cup_{n=1}^{\infty} I_n$ is a left (or right) ideal.

EXERCISE 9.5.14 Prove the left-sided case of Theorem 9.5.13.

EXERCISE 9.5.15 Let R be a commutative ring and Z the set of all zero divisors in R. What is wrong with the following proof that $Z \cup \{0\}$ is an ideal in R?

Proof. Suppose R is a commutative ring and Z is the set of the zero divisors in R.

(Y1) Let $z_1, z_2 \in Z \cup \{0\}$. If $z_1 = z_2 = 0$, then clearly $z_1 + z_2 = 0 \in Z \cup \{0\}$. If precisely one of z_1, z_2 is zero, then without loss of generality, $z_1 = 0$ and $z_2 \neq 0$. Then z_2 is a zero divisor, so there exists nonzero $a \in R$ such that $z_1 a = 0$. Thus $(z_1 + z_2)a = z_2 a = 0$, so that $z_1 + z_2$ is a zero divisor. If neither z_1 nor z_2 is zero, then there exist nonzero $a_1, a_2 \in R$ such that $z_1 a_1 = z_2 a_2 = 0$. Thus

$$(z_1 + z_2)a_1 a_2 = z_1 a_1 a_2 + z_2 a_2 a_1 = 0a_2 + 0a_1 = 0$$

so that $z_1 + z_2$ is a zero divisor. In any case $z_1 + z_2 \in Z \cup \{0\}$, so that $Z \cup \{0\}$ is closed under addition.

(Y2) By definition, $0 \in Z \cup \{0\}$.

(Y3) Let $z \in Z \cup \{0\}$. If $z = 0$, then $-z = 0 \in Z \cup \{0\}$. If $z \neq 0$, then there exists nonzero $a \in R$ such that $za = 0$. Thus $(-z)(a) = -za = -0 = 0$, so that $-z$ is a zero divisor. In either case $-z \in Z \cup \{0\}$.

(Y4) Let $z \in Z \cup \{0\}$, and $r \in R$. If $z = 0$, then $rz = 0 \in Z \cup \{0\}$. If $z \neq 0$, then there exists nonzero $a \in R$ such that $za = 0$. Since $a(rz) = r(az) = 0$, it follows that rz is a zero divisor.

Since $Z \cup \{0\}$ satisfies properties Y1–Y4, $Z \cup \{0\}$ is an ideal in R. □

EXERCISE 9.5.16 Demonstrate the error in the proof from the previous exercise by using appropriate elements from $\mathbb{Z}_{12}$.

9.6 Generated Ideals

Analogous to the subgroup generated by a subset of a group, we can define a similar term for ideals.

Definition 9.6.1 Suppose R is a ring and A is a non-empty subset of R. Suppose I is a subset of R with the following properties.

(U1) $A \subseteq I$.

(U2) I is a left ideal of R.

(U3) If J is a left ideal of R and $A \subseteq J$, then $I \subseteq J$.

Then I is called a *left ideal generated by* A and is denoted $(A)_l$.

We can similarly define a *right ideal generated by* A and denote it $(A)_r$. If R is commutative, we simply write (A) and call it the *ideal generated by* A. Does $(A)_l$ exist? If so, is it unique? And if so, what does it look like?

Exercise 9.6.2 Suppose R is a ring and A is a non-empty subset of R. Let $\mathcal{F}$ be the family of all left ideals of R that contain all elements of A. Then $(A)_l$ exists uniquely and can be written as

$$(A)_l = \bigcap_{I \in \mathcal{F}} I \tag{9.43}$$

If our path through groups and subgroups suggests a direction for us to go from here, we would consider the left ideal generated by a single ring element a and show that the top-down form of $(a)_l$ in Exercise 9.6.2 is equivalent to a form that can be built from the bottom up, by starting with a and building up to a subset of R that has properties U1–U3. Let's do that now. But in order for this program to be sure to work, R must have a unity element.

Theorem 9.6.3 Suppose R is a ring with unity and $a \in R$. Then $(a)_l = Ra$, the construction in Eq. (9.41), and $(a)_r = aR$.

Proof. We prove for $(a)_l$ only by showing that Ra satisfies properties U1–U3. The proof for $(a)_r$ is similar. First, since R has a unity e, $a = ea \in Ra$, so that U1 is satisfied. Second, by Exercise 9.5.5, Ra is an ideal of R, so that U2 is satisfied. Finally, suppose J is any left ideal of R that contains a, and pick any $x \in Ra$. Then $x = ra$ for some $r \in R$. But since J is an ideal of R and $a \in J$, it must be that $ra \in J$, so that $Ra \subseteq J$. Thus U3 is satisfied, and $(a)_l = Ra$. □

Example 9.6.4 Describe the left ideal of $\mathbb{Z}_{2\times 2}$ generated by $\begin{bmatrix} 2 & 0 \\ 0 & 3 \end{bmatrix}$.

Solution Let $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$ be any matrix in $\mathbb{Z}_{2\times 2}$. Then

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix}\begin{bmatrix} 2 & 0 \\ 0 & 3 \end{bmatrix} = \begin{bmatrix} 2a & 3b \\ 2c & 3d \end{bmatrix} \tag{9.44}$$

Thus the left ideal is the set of all 2 × 2 matrices whose first column entries are even and whose second column entries are multiples of three. ■

Exercise 9.6.5 Describe the right ideal of $\mathbb{Z}_{2\times 2}$ generated by the matrix in Example 9.6.4.

Regardless of whether the presence of a unity allows $(a)_l$ to be written in the form Ra, the left ideal generated by a single element $a \in R$ is called the *principal left ideal* generated by a. If R does not have a unity element, then Ra will not contain a. Thus the construction of (a) would be a bit more complicated than that in Theorem 9.6.3. In this text, (a) will always be in the context of a ring with unity.

Exercise 9.6.6 In the integers, what is $(6) \cap (15)$?

Exercise 9.6.7 In $\mathbb{Z}[t]$, describe (t).

Exercise 9.6.8 Find the principal left and right ideals generated by $\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$ in $\mathbb{Z}_{2\times 2}$.

Exercise 9.6.9 Let $M = \left\{ \begin{bmatrix} a & 0 \\ b & c \end{bmatrix} : a, b, c \in \mathbb{Z} \right\}$. Let $r = \begin{bmatrix} 2 & 0 \\ 6 & 3 \end{bmatrix}$, and consider the following subsets of M.

$$rM = \{rx : x \in M\} \tag{9.45}$$

$$Mr = \{xr : x \in M\} \tag{9.46}$$

(a) Show that M is a ring by showing it is a subring of $\mathbb{Z}_{2\times 2}$.

(b) By Exercise 9.5.5, the sets in Eqs. (9.45) and (9.46) are ideals in M. Show that neither is a subset of the other.

(c) Show that $rM \cap Mr$ is neither a left nor a right ideal in M.

Exercise 9.6.10 Suppose R is a ring with unity, and suppose a is a unit of R. Show that $(a)_l = (a)_r = R$.

We can construct a form analogous to Eq. (9.41) for the ideal generated by a finite set $\{a_1, a_2, \ldots, a_n\}$. To simplify the notation, we will denote this ideal $(a_1, \ldots, a_n)$.

Exercise 9.6.11 Suppose R is a ring with unity and $A = \{a_1, a_2, \ldots, a_n\}$ is a subset of R. Construct a form of $(A)_l$ that is analogous to Eq. (9.41) and show that it satisfies U1–U3.

Exercise 9.6.12 Show that the ideal defined in Exercise 9.5.3 is actually $(2, t)$.

Exercise 9.6.13 Let a and b be nonzero integers, and let $g = \gcd(a, b)$. Show that $(a, b) = (g)$.

Exercise 9.6.14 Let U be a non-empty set, and consider the ring on $\mathcal{P}(U)$ with symmetric difference and intersection (Exercise 9.1.8). Let A be a subset of U. Describe the ideal generated by A.

If R is a commutative ring with unity, there is an important link between the existence of proper ideals and the kind of ring R is. We will prove one direction of the next theorem, and you will prove the other.

Theorem 9.6.15 Suppose R is a commutative ring with unity. Then R is a field if and only if it has no proper ideals.

Proof.

($\Leftarrow$) Suppose R has no proper ideals. Choose any nonzero $x \in R$. We show that x has a multiplicative inverse. Since x is nonzero and R has no proper ideals, $(x) = R$, so that $e \in (x)$. By Theorem 9.6.3, $(x) = Rx$, so that there exists $r \in R$ such that $e = rx$. Thus x has a multiplicative inverse, and R is therefore a field. □

Exercise 9.6.16 Prove the $\Rightarrow$ direction of Theorem 9.6.15.

Suppose I is a proper ideal of a ring R and $a \in R - I$. Similar to the way we adjoin an element to a ring to create an extension, we can construct a left or right ideal of R that contains a and all elements of I. Our next result addresses the construction for the left-sided case. We will need this construction in Section 9.15 in the context of a commutative ring with unity.

Exercise 9.6.17 Let R be a ring, I an ideal of R, and fix $a \in R - I$. Let

$$J = \{ra + i : r \in R, i \in I\} \tag{9.47}$$

Then J is a left ideal of R.

9.7 Prime and Maximal Ideals

If p is a prime number and $p \mid ab$, then either $p \mid a$ or $p \mid b$ (Exercise 2.5.11). Another way to say this is if there exists integer k_1 such that $pk_1 = ab$, then there will exist integer k_2 such that either $pk_2 = a$ or $pk_2 = b$. Using the language of the ideal in the integers generated by p and Theorem 9.6.3, yet another way to say this is if $ab \in (p)$, then either $a \in (p)$ or $b \in (p)$. In a general ring, we assign a term to any *proper* ideal with this special property, whether or not the ideal is principal.

Definition 9.7.1 Suppose R is a ring, and P is a proper ideal of R with the property that $ab \in P$ implies either $a \in P$ or $b \in P$. Then P is called a *prime ideal* of R.

Exercise 9.7.2 Determine with proof whether each of the following ideals is prime.

(a) (6) and (7) in $\mathbb{Z}$

(b) $(2, t)$ in $\mathbb{Z}[t]$

(c) (t) in $\mathbb{Z}[t]$

(d) The left ideal from Exercise 9.6.8

Even though we have not yet looked into properties of prime elements in a ring, we must be careful about jumping to conclusions about prime ideals and principal ideals generated by prime elements. In Section 9.8, we will see that if a principal ideal is a prime ideal, then its generator must be a prime element. However, just because an element is prime, it does *not* mean that it generates a prime ideal. We will illustrate an example of how this can happen in the next section. In Section 9.10, we will see that a prime element in a more specialized kind of ring called a principal ideal domain will always generate a prime ideal.

Any proper ideal of a ring is contained in a larger ideal, namely, the ring itself. But if an ideal is as large as it can be without actually being the entire ring, then we call it *maximal*. Since we are primarily interested in maximal ideals in a commutative ring, we set the definition in that context.

Definition 9.7.3 Suppose R is a commutative ring, and I is a proper ideal with the property that if J is any ideal such that $I \subseteq J \subseteq R$, then either $J = I$ or $J = R$. Then I is called a *maximal* ideal.

One way to visualize a maximal ideal M is the following. If M is a maximal ideal, then for any $r \in R - M$, the only ideal of R that contains r and all elements of M is R itself.

Exercise 9.7.4 If p is a prime number, then the ideal it generates in the integers is maximal.

Exercise 9.7.5 Explain why (4) is not maximal in the integers.

Exercise 9.7.6 In $\mathbb{Z}_{2\times 2}$, the right ideal $I = \left\{ \begin{bmatrix} 2a & 2b \\ 2c & 2d \end{bmatrix} : a, b, c, d \in \mathbb{Z} \right\}$ is not maximal, for $J = \left\{ \begin{bmatrix} 2a & 2b \\ c & d \end{bmatrix} : a, b, c, d \in \mathbb{Z} \right\}$ is also a right ideal.

Exercise 9.7.7 In the power set ring (Exercise 9.1.8), let $A_0 \subseteq U$. Then $P(U - A_0)$ is a maximal ideal if and only if A_0 contains precisely one element.

We see immediately that there is a relationship between prime and maximal ideals in commutative rings with unity.

Exercise 9.7.8 If M is a maximal ideal in a commutative ring with unity, then M is prime.

The reason that the ring in Exercise 9.7.8 has to have a unity element can be seen in the next exercise.

Exercise 9.7.9 Show that (4) is maximal but not prime in the ring of even integers.

About now you should be asking either for a theorem claiming that prime ideals in commutative rings with unity are also maximal (so that the terms are logically equivalent) or for an example of a prime ideal that is not maximal. Well, they are not logically equivalent.

Exercise 9.7.10 Demonstrate an example of a prime ideal that is not maximal.[4]

When we restrict ourselves to principal ideal domains in Section 9.10, we will see that prime ideals are also maximal.

Exercise 9.7.11 In $\mathbb{Z} \times \mathbb{Z}$, is $(5) \times (3)$ prime and/or maximal? Verify your answer.

Exercise 9.7.12 For non-empty set U and $A_0 \subseteq U$, show that the ideal $P(U - A_0)$ is prime if and only if A_0 has precisely one element.

9.8 Integral Domains

A commutative ring with unity that has no zero divisors is called an *integral domain*, or *domain* for short. The absence of zero divisors means that the principle of zero products applies by definition. Thus the integers are a standard first example of a domain. Since all nonzero elements of a field are units and a unit is not a zero divisor (Exercise 9.3.16), a field contains no zero divisors and is therefore a domain. Thus the rationals, reals, and complex numbers are all domains.

Exercise 9.8.1 If p is a prime number, then $\mathbb{Z}_p$ is a domain.[5]

Example 9.8.2 $\mathbb{Z}[\sqrt[3]{2}]$ and $\mathbb{Q}_{OD}$ are domains, for as subrings of the real numbers, they are commutative, contain 1, and have no zero divisors. ■

Example 9.8.3 If U is a set with more than one element, then the power set ring from Exercise 9.1.8 is not an integral domain, for every non-empty, proper subset of U is a zero divisor (Exercise 9.3.15). ■

Exercise 9.8.4 If D is a domain, and if $ac = bc$ and $c \neq 0$, then $a = b$.

[4] You shouldn't have to look far.

[5] See Exercise 2.5.11.

Exercise 9.8.5 Suppose D is a domain and a and b are domain elements such that $a \mid b$. Then the element k such that $ak = b$ is unique.

With Exercise 9.8.5, the following term becomes meaningful.

Definition 9.8.6 Suppose D is a domain, $a, b \in D$, and a is not a unit. If $a \mid b$, where $ak = b$ and k is not a unit, then a is called a *proper* divisor of b.

In a ring in which the result of Exercise 9.8.5 does not apply, such as those in Exercise 9.3.10, then Definition 9.8.6 cannot be unambiguously applied. For it is possible to have $ak_1 = b$ and $ak_2 = b$ in a ring R where k_1 is a unit in R but k_2 is not.

Example 9.8.7 In $\mathbb{Z}_{2\times 2}$,

$$\begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 2 \\ 1 & 2 \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}\begin{bmatrix} 2 & 1 \\ -1 & 1 \end{bmatrix} \tag{9.48}$$

Now $\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$ is a unit in $\mathbb{Z}_{2\times 2}$, for

$$\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}\begin{bmatrix} 1 & -1 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & -1 \\ 0 & 1 \end{bmatrix}\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \tag{9.49}$$

while $\begin{bmatrix} 2 & 1 \\ -1 & 1 \end{bmatrix}$ is not a unit in $\mathbb{Z}_{2\times 2}$. ■

Exercise 9.8.8 Show that $\begin{bmatrix} 2 & 1 \\ -1 & 1 \end{bmatrix}$ is not a unit in $\mathbb{Z}_{2\times 2}$.

Multiplicative cancellation in a domain is a logical consequence of the principle of zero products. We could have defined a domain as a commutative ring with unity where multiplicative cancellation holds for nonzero elements, and we could have then shown that the principle of zero products follows from that. In a commutative ring, the principle of zero products and multiplicative cancellation are logically equivalent.

Exercise 9.8.9 A commutative ring with unity and multiplicative cancellation is a domain.

If a commutative ring with unity has a nonzero characteristic, then the only way it can be a domain is if the characteristic is prime.

Exercise 9.8.10 If a domain has nonzero characteristic n, then n is prime.

With Theorem 9.3.25 and Exercises 9.8.1 and 9.8.10, we see that $\mathbb{Z}_n$ is a domain if and only if n is prime. The only feature a field has that a domain might not is the existence of multiplicative inverses. Thus there is only one thing to show in the next exercise.

EXERCISE 9.8.11 A finite integral domain is a field.[6]

By Exercise 9.8.11, we have that $\mathbb{Z}_p$ is a field if and only if p is prime.

EXERCISE 9.8.12 Find multiplicative inverses for all nonzero elements of $\mathbb{Z}_7$.

In the integers, you showed in Exercise 2.5.5 that if $a \mid b$ and $b \mid a$, then $a = \pm b$. Furthermore, ± 1 are the units in the integers. In a general domain, we say that a and b are *associates* if $a \mid b$ and $b \mid a$. Since this definition does not apply to zero, we declare zero to be an associate of itself. Notice that this declaration and the definition of associate prevent zero from having any other associates in a domain.

EXERCISE 9.8.13 Show that $1 + i$ and $1 - i$ are associates in the Gaussian integers $\mathbb{Z}[i]$.

EXERCISE 9.8.14 In a domain, two elements a and b are associates if and only if there exist units u and v such that $a = ub$ and $b = va$.

You should feel an equivalence relation coming on about now.

EXERCISE 9.8.15 Let D be a domain, and define $a \sim b$ if a is an associate of b. Then $\sim$ is an equivalence relation on D.[7]

Anytime you create an equivalence relation, it is natural to ask two questions: What do the equivalence classes look like, and which element from each equivalence class might be a good choice as a representative element of the class?

EXERCISE 9.8.16 If D is a domain and $\sim$ is the equivalence relation of association, then $[e]$ is the set of units of D.

If a and b are associates, then there ought to be some senses in which they are interchangeable. One example of how this is true is in the integers, where $(6) = (-6)$. Associates generate the same principal ideal. The best way to show this is first to prove the following. Its corollary is immediate.

EXERCISE 9.8.17 Suppose a and b are elements of a domain. Then $a \mid b$ if and only if $(b) \subseteq (a)$.

[6] To find a^{-1}, define $f(x) = ax$ and apply Exercise 4.6.11.

[7] Don't forget zero.

Corollary 9.8.18 In a domain, $(a) = (b)$ if and only if a and b are associates.

If a is a proper divisor of b, then by Exercises 9.8.5 and 9.8.14, a and b are not associates. So $(b) \subseteq (a)$ by Exercise 9.8.17, but $(a) \neq (b)$ by Corollary 9.8.18. Thus we have the following.

Corollary 9.8.19 If a is a proper divisor of b in a domain, then (b) is a proper subset of (a).

Now let's look at the relationship between prime elements and prime ideals.

EXERCISE 9.8.20 If D is a domain and (p) is a prime ideal, then p is prime in D.

But what about the converse? If p is a prime element of a domain, must it generate a prime ideal? In the integers the answer is yes. Thanks to Exercise 2.5.11, if p is a prime number, then (p) is a prime ideal in the integers. But in the integers, Exercise 2.5.11 depends on the existence of greatest common divisors, which itself depends on the WOP. Strangely, in some domains a prime element can generate an ideal that is not prime, and a prime element p might satisfy $p \mid ab$, while p divides neither a nor b. The next exercise illustrates these possibilities. We provide the solution to the first part and leave the rest to you. We will return to this example in Section 9.9.

EXERCISE 9.8.21 The ring $\mathbb{Z}[\sqrt{-5}]$ is a domain because it is a subring of the complex numbers and contains 1. However, you can show the following.

(a) 3 is prime in $\mathbb{Z}[\sqrt{-5}]$.

Solution First note

$$9 = 3 \cdot 3 = (2 + \sqrt{-5})(2 - \sqrt{-5}) \tag{9.50}$$

To show 3 is prime in $\mathbb{Z}[\sqrt{-5}]$, suppose

$$3 = (a + b\sqrt{-5})(c + d\sqrt{-5}) \tag{9.51}$$

where a, b, c, d are integers. We show that one of these two factors must be a unit, which by Theorem 9.4.3 must be ± 1. Multiplying out the factors and using the definition of equality in $\mathbb{Z}[\sqrt{-5}]$ yields

$$ac - 5bd = 3 \quad \text{and} \quad ad + bc = 0 \tag{9.52}$$

Proceeding as in the proof of Theorem 9.4.3, we square each equation, multiply the latter by 5, add and factor to have

$$(a^2 + 5b^2)(c^2 + 5d^2) = 9 \tag{9.53}$$

Since both factors in Eq. (9.53) are positive integers, both must be in the set $\{1, 3, 9\}$. Considering each possibility reveals either a contradiction or the fact that one of the factors is a unit. Thus 3 is prime in $\mathbb{Z}[\sqrt{-5}]$.

(b) $2 \pm \sqrt{-5}$ are prime in $\mathbb{Z}[\sqrt{-5}]$.

(c) $2 \pm \sqrt{-5} \notin (3)$.

(d) There exists a prime element $p \in \mathbb{Z}[\sqrt{-5}]$ such that (p) is not a prime ideal.

(e) There exist $a, b, p \in \mathbb{Z}[\sqrt{-5}]$ where p is prime, $p \mid ab$, but p divides neither a nor b.

In the polynomial ring over a domain, the degree of a product of polynomials behaves in a more predictable way than it might in the polynomial ring over an arbitrary ring.

EXERCISE 9.8.22 Suppose D is a domain and $f, g \in D[t]$ are nonzero polynomials. Then

(a) $\deg(fg) = \deg f + \deg g$.

(b) If $f \mid g$, then $\deg f \leq \deg g$.

EXERCISE 9.8.23 The polynomial ring over a domain is a domain.

The next exercise will help you see more clearly what it means for a polynomial in $D[t]$ to be prime.

EXERCISE 9.8.24 Let D be a domain whose set of units is U, and let K be a field.

(a) What are the units of $D[t]$?

(b) What are the units in $K[t]$?

In a polynomial ring, we generally use the word *irreducible* instead of prime to refer to such a polynomial.

EXERCISE 9.8.25 Show that $f = t^2 - 2$ is irreducible in $\mathbb{Z}[t]$. Let $\mathbb{Z}[\sqrt{2}, t]$ represent the polynomial ring over the domain $\mathbb{Z}[\sqrt{2}]$. Show that f is reducible in $\mathbb{Z}[\sqrt{2}, t]$.

EXERCISE 9.8.26 Is $f = 2t^2 - 4$ reducible in $\mathbb{Z}[t]$? In $\mathbb{Q}[t]$?

If $f = a_n t^n + \cdots + a_0$ is a polynomial in $\mathbb{Z}[t]$ such that $\deg f \geq 1$, we call $\gcd(a_n, a_{n-1}, \ldots, a_0)$ the *content* of f. If the content of f is one, we say that f is *primitive*. Insight from Exercise 9.8.26 should make your proof of the following immediate.

Exercise 9.8.27 Suppose $f \in \mathbb{Z}[t]$ is irreducible and $\deg f \geq 1$. Then f is primitive.

Exercises 9.8.26 and 9.8.24 reveal that a result similar to Exercise 9.8.27 does not apply to $\mathbb{Q}[t]$. That is, if a polynomial with integer coefficients is irreducible when viewed as an element $\mathbb{Z}[t]$, then it is primitive. However, $2t + 6$ is irreducible in $\mathbb{Q}[t]$ but has content 2.

Exercise 9.8.28 Suppose R is a ring and let r be a nonzero element. Define $f : R \to R$ by $f(x) = rx$. Show that f is not necessarily one-to-one. What additional condition on R guarantees f is one-to-one? Prove.

Exercise 9.8.29 Suppose R and S are both domains. Does it follow that $R \times S$ is a domain?

9.9 Unique Factorization Domains

Theorem 3.5.19 states that every integer greater than 1 has a unique factorization into positive integer primes. If we broaden the context to the integers, we lose uniqueness in most cases because we could introduce some negative signs here and there. But this is the only way we lose uniqueness. Even then, the prime factors in two different factorizations can be paired up as associates of each other (additive inverses), allowing us to say that every integer greater than 1 has a prime factorization in the integers that is unique up to order and association of the factors. For a negative integer n, applying the same principle to $-n$ allows us to say that every nonzero integer that is not a unit has a factorization into prime integers, and this factorization is unique up to order and association of the factors. An integral domain in which every nonzero element has a factorization into prime elements, unique up to order and association of the factors, is called a *unique factorization domain*, or UFD for short. In a field every nonzero element is a unit, so a field is trivially a UFD.

In our progression from more general to more specialized rings, UFDs are the point where we can show that any two nonzero elements have a greatest common divisor, unique up to association. We will adapt Definition 2.5.6 slightly to make it applicable to a general domain and then take a quick look at how we show existence and some sort of uniqueness of the greatest common divisor in a UFD. The first place we actually need the existence of the gcd is in a principal ideal domain (PID), where it is pretty easy to show.

An important example of a UFD that we merely make reference to right now is $\mathbb{Z}[t]$. To verify that $\mathbb{Z}[t]$ is a UFD takes some mathematical machinery that we will create in Section 9.11, where we will study $\mathbb{Q}[t]$ in some depth. Even though $\mathbb{Z}[t]$ is a subring of $\mathbb{Q}[t]$, it is not as specialized a ring because there are limitations on coefficients in $\mathbb{Z}[t]$ that do not apply in $\mathbb{Q}[t]$. To show that $\mathbb{Z}[t]$ is a UFD, it helps to view it in the context of $\mathbb{Q}[t]$.

Since a UFD is by definition an integral domain, we do not need a theorem claiming that a UFD is a domain. But not all domains are UFDs. For example, $\mathbb{Z}[\sqrt{-5}]$ is a domain, and Eq. (9.50) shows that 9 has two distinct factorizations into prime elements. Thus $\mathbb{Z}[\sqrt{-5}]$ is not a UFD.

Definition 9.9.1 Suppose D is a domain, and let a and b be nonzero elements of D. Suppose g is a domain element with the following characteristics:

(D1) $g \mid a$ and $g \mid b$.

(D2) If h is any element of D with the properties that $h \mid a$ and $h \mid b$, then it must be that $h \mid g$ also.

Then g is called a *greatest common divisor* of a and b and is denoted $\gcd(a, b)$.

In Section 2.5, we said that a practical way you might find $\gcd(a, b)$ in the positive integers is by breaking down a and b into their prime factorizations, taking the appropriate number of 2s, 3s, and so on, and building the gcd from that. We did not prove that such a trick produces a positive integer with properties D1–D2 because the WOP provided an easier and more useful way. Alas, in a general UFD, the existence of $\gcd(a, b)$, unique up to association, must be proved by exploiting the unique prime factorizations of a and b in that somewhat sloppy way. We state the theorem here, followed by some details of the proof that might make the notation minimally sloppy. You will write the complete proof as an exercise.

Theorem 9.9.2 Suppose D is a UFD and a and b are nonzero elements of D. Then there exists $g \in D$ that satisfies D1–D2, and if g_1 and g_2 both satisfy D1–D2, then g_1 and g_2 are associates.

Here are some suggestions on how to prove Theorem 9.9.2 and keep the notation from getting outrageously complicated. If we break a and b down into prime factorizations, then let $\{p_1, p_2, \ldots, p_n\}$ be all the primes that appear in either a or b, we can write

$$a = p_1^{\alpha_1} p_2^{\alpha_2} \ldots p_n^{\alpha_n} \quad \text{and} \quad b = p_1^{\beta_1} p_2^{\beta_2} \ldots p_n^{\beta_n} \tag{9.54}$$

where some of the α_k and β_k might be zero. If we let $\gamma_k = \min\{\alpha_k, \beta_k\}$ for all $1 \le k \le n$, you can show that $g = p_1^{\gamma_1} p_2^{\gamma_2} \ldots p_n^{\gamma_n}$ satisfies D1–D2. Since $\gamma_k = \min\{\alpha_k, \beta_k\}$, we know that $\gamma_k \le \alpha_k$ and $\gamma_k \le \beta_k$ for all k. This should make it easy to show that g has property D1. To show that g has property D2, suppose $h \mid a$ and $h \mid b$. Then there exist $k_1, k_2 \in D$ such that $hk_1 = a$ and $hk_2 = b$. If the unique factorization of h is written as $h = q_1^{\delta_1} q_2^{\delta_2} \ldots q_m^{\delta_m}$, then the prime factorizations of hk_1 and hk_2 must agree with Eqs. (9.54), so that some possible reordering of the p_k

allows us to say that $q_k = p_k$ for all $1 \leq k \leq m \leq n$. If $m < n$, then letting $\delta_k = 0$ for $m + 1 \leq k \leq n$ gives us $h = p_1^{\delta_1} \dots p_n^{\delta_n}$. But the fact that $h \mid a$ implies that $\delta_k \leq \alpha_k$ for all $1 \leq k \leq n$. Similarly, $h \mid b$ implies that $\delta_k \leq \beta_k$ for all $1 \leq k \leq n$. Thus $\delta_k \leq \gamma_k$ for all $1 \leq k \leq n$, and by constructing the appropriate element of D, you can show $h \mid g$. Thus g has property D2. Showing g is unique up to association is surprisingly easy. In fact, the way you proved uniqueness of $\gcd(a, b)$ in the positive should translate directly over to D to imply that any g_1 and g_2 that satisfy D1–D2 are associates.

EXERCISE 9.9.3 Prove Theorem 9.9.2.

9.10 Principal Ideal Domains

Since an integral domain is commutative and has a unity element, left and right ideals are indistinguishable, and the principal ideal generated by an element a can always be written as Da. Furthermore, by Exercise 9.6.11, if $A = \{a_1, \dots, a_n\}$, then

$$(A) = \{d_1a_1 + d_2a_2 + \cdots + d_na_n : d_k \in D \text{ for all } 1 \leq k \leq n\} \tag{9.55}$$

Let's write this form of the ideal generated by A as $Da_1 + Da_2 + \cdots + Da_n$.

In some domains, principal ideals are the only ones there are. Every ideal will have a single generator. If a domain is such that every ideal is principal, it is called a *principal ideal domain*, or a PID for short. A field is trivially a PID, for Exercise 9.6.16 says that a field has no proper ideals. The trivial ideal in a field is generated by zero, and the field itself can be generated by any nonzero element.

You have already seen an example of an integral domain that is not a PID, but before we point it out, let's see why some of the domains we already know are PIDs, and derive some results for PIDs in general. Then the example of a domain (actually a UFD) that is not a PID will point itself out. First let's review the integers a bit.

In Chapter 2, after we defined divisibility and gcd in the integers, we showed that the gcd exists uniquely for all nonzero integer pairs (Theorem 2.5.7). Furthermore, it is the smallest positive element of

$$S = \{ma + nb : m, n \in \mathbb{Z}\} \tag{9.56}$$

Thus our proof that gcds exist in the integers depended on the WOP to give us a number that we could show satisfies properties D1–D2. Furthermore, in showing that this smallest element of S divides both a and b, you employed the division algorithm. And you proved the division algorithm with the help of the WOP.

Notice that S in Eq. (9.56) is precisely $\mathbb{Z}a + \mathbb{Z}b$, the ideal generated by $\{a, b\}$. So $\gcd(a, b)$ is the smallest element of (a, b). Moreover, in Exercise 9.6.13, you showed that $\{a, b\}$ and $\gcd(a, b)$ generate the same ideal in the integers, so that the ideal generated by a two-element set is really principal after all. Once again, it is the WOP at work, leading you to the smallest positive element of an ideal, and allowing you to show that it generates the whole ideal. By looking in this same direction, you can show that the integers are a PID. The way you will attack this in the next exercise deserves a comment.

To show that any particular domain D is a PID, you must show that every ideal can be written as Da, where a is a generator of the ideal that you must find. For the trivial ideal, that's a piece of cake; otherwise, the ideal has a positive element. So the WOP comes in and gives you a generator on a silver platter (details notwithstanding). Thus the WOP is the basis for the integers being a PID and for its containing gcds. But the WOP depends on $<$ as a measure of size of elements in the positive integers.

EXERCISE 9.10.1 The integers are a PID.

Another example of a PID is $\mathbb{Q}[t]$. Rather than show this directly, we will show in Section 9.11 that $\mathbb{Q}[t]$ is a Euclidean domain and that all Euclidean domains are PIDs.

A creative use of the WOP can supply you with a generator to prove the following.

EXERCISE 9.10.2 Show that $\mathbb{Q}_{OD}$ is a PID. See the hint if you need help finding a generator.[8]

Now suppose we're working in a domain in which we don't necessarily have a measure of element size, so that we have no way to apply the WOP. With the assumption that all ideals are principal, you can prove that gcds exist in a PID.

EXERCISE 9.10.3 Suppose D is a PID and a and b are nonzero elements of D. Then $\gcd(a, b)$ exists.[9]

Since a generator for (a, b) qualifies as $\gcd(a, b)$ and since $(a, b) = Da + Db$, we have the following.

Corollary 9.10.4 If D is a PID and a and b are nonzero elements of D, then there exists $m, n \in D$ such that $\gcd(a, b) = ma + nb$.

Unfortunately, Exercise 9.10.3 says nothing about the uniqueness of $\gcd(a, b)$. But there is something close.

EXERCISE 9.10.5 If g_1 and g_2 are both greatest common divisors of a and b in a PID, then g_1 and g_2 are associates.

[8] Of all nonzero elements of an ideal, pick one with a smallest power of 2 in the numerator.

[9] Let g be a generator of the ideal generated by a and b.

So even if there are several gcds of a and b, they are all unit multiples of each other. If g is one gcd of a and b, then everything in its equivalence class of associates is, too. If $[e]$ is the set of gcds of a and b, then we say a and b are *relatively prime*, and we note that if a and b are relatively prime, then there exists a linear combination such that $ma + nb = e$.

Now let's find that example of a domain (actually, a UFD) that is not a PID. In any commutative ring, a maximal ideal is prime (Exercise 9.7.8). But some prime ideals are not maximal, such as (t) in $\mathbb{Z}[t]$ (Exercise 9.7.10). And (t) is a proper subset of $(2, t)$, the ideal from Exercise 9.5.3 of all polynomials with even constant term. The following result lets us close in for the kill.

EXERCISE 9.10.6 A prime ideal in a PID is maximal.

Corollary 9.10.7 $\mathbb{Z}[t]$ is not a PID.

Proof. $(2, t)$ is prime but not maximal in $\mathbb{Z}[t]$. □

In Section 9.11, we will show that $\mathbb{Z}[t]$ is a UFD.

EXERCISE 9.10.8 Show that $(2, t)$ is not a principal ideal of $\mathbb{Z}[t]$ by showing that every element fails to generate the ideal.

Now we want to show that a PID is a UFD. It takes several steps to get there, but fortunately a lot of the uniqueness work has already been done for the positive integers in the proof of Theorem 3.5.19 and translates over to a PID almost word for word. The sticky part in showing the existence of prime factorizations in a PID stems from the fact that the defining characteristic of a PID concerns its ideals, while the defining characteristic of a UFD concerns its elements. The link between them is the following. All ideals in a PID have a generator, and the way these ideals contain each other as subsets is tied by Exercise 9.8.17 to divisibility of their generators. Since in a PID $a \mid b$ if and only if $(a) \supseteq (b)$, the way an element breaks down into factors is directly linked to the way principal ideals stack. So let's make an important observation about the way ideals in a PID *cannot* stack.

Let D be a PID and $\{I_n\}_{n=1}^{\infty}$ a family of ideals such that $I_n \subseteq I_{n+1}$ for all n. By Theorem 9.5.13, $I = \cup_{n=1}^{\infty} I_n$ is an ideal. Since D is a PID, I is principal and has a generator a. Now since $a \in \cup_{n=1}^{\infty} I_n$, a is in some I_n. We claim that for all $k \geq n$, $I_k = I_n$. Suppose that $k \geq n$, and let x be any element of I_k. Then $x \in I$, so that $x = ay$ for some $y \in D$. But since I_n is an ideal that contains a, then $ay \in I_n$. Thus $I_k \subseteq I_n$, so that $I_k = I_n$.

The upshot of all this is that if you are in a PID and have a set of hypothesis conditions that allows you to create a family of ideals $\{I_n\}_{n=1}^{\infty}$ where I_n is a proper subsest of I_{n+1} for all n, then you have produced a contradiction, and at least some of the hypothesis conditions are false in a PID.

Now let's show that every element of a PID can be written as a product of prime elements and that this factorization is unique up to order and association of the factors. To prove the former, suppose a PID contains some element a that cannot be written as a product of primes. Then a is not prime itself. So there exist domain elements a_1 and b_1 such that $a = a_1b_1$ and neither a_1 nor b_1 is a unit. Furthermore, since a cannot be written as a product of primes, then a_1 or b_1 (or both) cannot either. Without loss of generality, we may assume a_1 is the culprit, which means that a_1 is not prime. And notice, since b_1 is not a unit, (a) is a proper subset of (a_1) by Corollary 9.8.19.

Now since a_1 is not prime, there exist a_2 and b_2 such that $a_1 = a_2b_2$ and neither a_2 nor b_2 is a unit. Once again, since a_1 cannot be written as a product of primes, either a_2 or b_2 cannot either. Assume it's a_2, and note that (a_1) is a proper subset of (a_2) because b_2 is not a unit. Continuing in the same way, we can generate an infinite set of ideals $\{(a_1), (a_2), \dots\}$ where each (a_n) is a proper subset of (a_{n+1}). This is a contradiction, so the assumption that there exists an element that cannot be written as a product of primes is false, and all elements have a factorization into primes. We have therefore proved the following.

Theorem 9.10.9 Every nonzero, non-unit element of a PID can be written as a product of primes.

To show the uniqueness of the prime factorization, almost all the work translates directly over from our work in the integers. The first step is analogous to Theorem 2.5.10, but the proof comes out a little differently because primes in a PID are defined in language that differs from that in the positive integers.

Theorem 9.10.10 If D is a PID, a is nonzero, and p is prime, then either $p \mid a$ or p and a are relatively prime.

Proof. Let $g = \gcd(a, p)$. Since $g \mid p$ and p is prime, then either g is a unit or it is an associate of p. If g is a unit, then writing $am + np = g$ and multiplying both sides through by g^{-1} reveals that a and p are relatively prime. If g is an associate of p, then $gu = p$ for some unit u, so that $g = pu^{-1}$. Also, the fact that $g \mid a$ means that $pu^{-1} \mid a$, so that $p \mid a$. □

With Theorem 9.10.10 in hand, the proofs of the next two theorems would be identical to those for Exercises 2.5.11 and 3.4.19.

Theorem 9.10.11 If D is a PID and $a, b, p \in D$ are such that p is prime, a and b are nonzero and $p \mid ab$, then either $p \mid a$ or $p \mid b$.

Theorem 9.10.12 Suppose D is a PID, and $p, a_1, a_2, \dots, a_n \in D$ are such that p is prime and $a_k \neq 0$ for all $1 \leq k \leq n$. Then if $p \mid a_1a_2 \dots a_n$, then there exists $k (1 \leq k \leq n)$ such that $p \mid a_k$.

The uniqueness result is analogous to Theorem 3.5.19, but the proof does not come out exactly the same because of the possible association of the factors. The next exercise will come in handy when you write the proof in the exercise that follows.

Exercise 9.10.13 If D is a domain, $p \in D$ is prime, and $u \in D$ is a unit, then pu is also prime.[10]

Exercise 9.10.14 If D is a PID and a is nonzero and not a unit, then the prime factorization from Theorem 9.10.9 is unique up to order and association of the factors.

With Theorem 9.10.9 and Exercise 9.10.14, we have the following.

Theorem 9.10.15 A PID is a UFD.

Finally, as a direct result of Theorem 9.10.11, we have the following.

Exercise 9.10.16 If D is PID and p is a prime element, then (p) is a prime ideal.

9.11 Euclidean Domains

Let's return to the division algorithm (Theorem 2.4.7) to expand it and cast it in somewhat different language. In writing $b = aq + r$, we insisted that a be positive, so that $0 \leq r < a$ is meaningful. We did not have to make this restriction on a. Obviously, a must be nonzero, but we could have allowed a to be negative to have a slightly stronger theorem. It would look something like this.

Theorem 9.11.1 (Extended Division Algorithm). Suppose a and b are nonzero integers. Then there exist unique integers q and r such that $b = aq + r$ and $0 \leq r < |a|$.

Proving to existence if a is negative is easy by appealing to existence for $-a$. In Theorem 9.11.1, we are using absolute value as a measure of the size of integers and are saying that any pair of nonzero integers can be related by breaking one of them (b) down into a certain multiple (q) of the other (a), with some possible stuff left over (r), but where the stuff left over is smaller in size than a. There are other important integral domains in which a notion of size can be imposed on the elements, and something akin to the division algorithm using that measure of size works for all nonzero elements. Since we do not need to apply the division algorithm to the zero element, we do not insist that zero be assigned a measure of size.

[10] If $pu = ab$ is a factorization of pu, then $a(bu^{-1})$ is a factorization of p.

Definition 9.11.2 Suppose D is an integral domain, and suppose there exists a function $d : D^{\times} \to \mathbb{W}$ with the property that $d(a) \leq d(ab)$ for all nonzero a and b. Such a function is called a *valuation*. Suppose also that for all nonzero elements a and b, there exist $q, r \in D$, such that $b = aq + r$, and either $r = 0$ or $d(r) < d(a)$. Then D is called a *Euclidean domain*.

Demonstrating that a domain is Euclidean requires the creation of a valuation that assigns a nonnegative integer size to all nonzero elements of the domain. There might be more than one such valuation, but below we will point out some features a valuation must have as a result of the requirement that $d(a) \leq d(ab)$. Also, given two nonzero elements a and b, b must either be a multiple of a, in which case $r = 0$, or be just the right distance from a multiple of a so that $b = aq + r$ can be written with r of sufficiently small size.

Example 9.11.3 The integers are Euclidean, since absolute value may serve as a valuation in Theorem 9.11.1. ■

Example 9.11.4 A field is a Euclidean domain. Let $d(a) = 1$ for all nonzero a, and note that $d(a) = d(ab)$ for all nonzero a and b. Also, we may let $q = ba^{-1}$ to have $b = aq$, so that $r = 0$ is always possible. ■

Divisibility in a field is trivial since nonzero elements are all multiples of each other. Exercise 9.11.7 will reveal that any valuation on a field must be constant.

Example 9.11.5 The Gaussian integers $\mathbb{Z}[i]$ are Euclidean, and the function defined by $d(a + bi) = a^2 + b^2$ can be shown to be a valuation that works. We will not provide a proof here. ■

Before we spend some quality time with one more very important example of a Euclidean domain, let's derive some results about Euclidean domains in general. First, in the next exercise you will show that a Euclidean domain is a PID. If D is Euclidean and I is an ideal, then you must find some ideal element of which every element of I is a multiple. But that should not be too hard, for an element whose valuation is minimal will serve nicely as a generator, and the division algorithm on $D^{\times}$ ought to be just the tool to enable you to show it.

EXERCISE 9.11.6 A Euclidean domain is a PID.

If D is Euclidean, it is by definition an integral domain. Thus it has a unity element e, and for any $a \in D$, it follows that $d(e) \leq d(ea) = d(a)$, so that $d(e)$ is minimal among all values of d. This fact and the division algorithm should come in handy when you prove the ⇒ direction of the following.

EXERCISE 9.11.7 In a Euclidean domain with valuation d, $d(u) = d(e)$ if and only if u is a unit.

Exercise 9.11.7 sheds a little more light on some things we already know. First, since absolute value can serve as a valuation on the integers, we see that the units in the integers are precisely the values of x for which $|x| = |1|$, namely, ± 1. Also, since every nonzero element of a field is a unit, any valuation d must satisfy $d(x) = d(e)$ for all nonzero x. Thus d must be constant. Conversely, if a valuation on a Euclidean domain is constant, then every nonzero element is a unit, so that the Euclidean domain is a field.

As we have progressed from rings to domains to UFDs to PIDs to Euclidean domains, we have claimed that there exist examples of one structure that do not qualify as an example of the next most restrictive structure. In every case up to now, we have provided an example complete with proof, except for $\mathbb{Z}[t]$. We have shown that $\mathbb{Z}[t]$ is not a PID (Corollary 9.10.7), but we have not yet shown it is a UFD. We will do this in the next section.

About now you should be asking, "Where is my example of a PID that's not Euclidean?" Well, this is the only place in our progression from more general to more specialized rings where we are going to present an example of a PID that is not Euclidean and will give only a loose explanation of how this would be shown. The classic example of a non-Euclidean PID was first constructed by T. Motzkin in 1949—very recently indeed by mathematical standards. It is $\mathbb{Z}[(1+\sqrt{-19})/2]$, the set of all expressions of the form $m + n(1+\sqrt{-19})/2$, where m and n are integers. For convenience, let's write $\alpha = (1+\sqrt{-19})/2$, and discuss how one goes about showing $\mathbb{Z}[\alpha]$ is a PID that is not Euclidean.

First, we address the fact that $\mathbb{Z}[\alpha]$ is not Euclidean. We would do this by showing that every Euclidean domain has a certain feature, then showing that $\mathbb{Z}[\alpha]$ does not have this feature. If D is any Euclidean domain that is not a field, then there will exist nonzero, non-unit elements. If d is a valuation, then a nonzero, non-unit element x will satisfy $d(x) > d(e)$. Among all elements of D, let a be a nonzero, non-unit element for which $d(a)$ is minimal. Such an element is called a *universal side divisor*. By the division algorithm on D, any nonzero x can be written as $x = aq + r$, where either $r = 0$ or $d(r) < d(a)$. Since $d(a)$ is minimal among all nonzero, non-unit elements of D, we may say that r is either zero or a unit. Every Euclidean domain that is not a field will have universal side divisors because the valuation will not be constant.

Next we need to know what the units are in $\mathbb{Z}[\alpha]$. It turns out that the only units in $\mathbb{Z}[\alpha]$ are ± 1. So let's suppose $\mathbb{Z}[\alpha]$ is a Euclidean domain and see how this leads to a contradiction. Since the only units in $\mathbb{Z}[\alpha]$ are ± 1, $\mathbb{Z}[\alpha]$ is not a field. Thus there exists a universal side divisor $a \in \mathbb{Z}[\alpha]$, and every $x \in \mathbb{Z}[\alpha]$ can be written as $x = aq + r$, where $r \in \{0, \pm 1\}$. In particular, $x = 2$ must be writable in this way. So $aq = 2 - r$, and the only possible values of $2 - r$ are $\{1, 2, 3\}$. Thus a divides at least one of $\{1, 2, 3\}$, but since a is not a unit, a does not divide 1. With some work, we could show that the only divisors of 2 are $\{\pm 1, \pm 2\}$ and the only divisors of 3 are $\{\pm 1, \pm 3\}$. Thus $a \in \{\pm 2, \pm 3\}$. But letting $x = \alpha$, it turns out that there is no $q \in \mathbb{Z}[\alpha]$ for which $\alpha = aq + r$, given that $a \in \{\pm 2, \pm 3\}$ and $r \in \{0, \pm 1\}$. This is a contradiction, so $\mathbb{Z}[\alpha]$ is not a Euclidean domain.

Now we address the fact that $\mathbb{Z}[\alpha]$ is a PID. First, let's take the defining characteristic of a Euclidean domain and state it first in the original way and then in an altered, but equivalent, form. A domain D is a Euclidean domain if there exists a valuation d on $D^{\times}$ such that for all nonzero a and b,

1. There exist $q, r \in D$ such that $b = aq + r$ and either $r = 0$ or $d(r) < d(a)$.
2. Either b is in the ideal generated by a, or it is possible to subtract from b some q multiple of a to produce an element $r = b - aq$ whose valuation is smaller than that of a, that is, $d(r) < d(a)$.

Let's weaken this second form a bit. Suppose D is a domain with a valuation such that for all nonzero a and b, either b is in the ideal generated by a, or there is some linear combination of a and b whose valuation is smaller than that of a. That is, either $b \in (a)$ or there exist nonzero domain elements m and n such that $d(ma + nb) < d(a)$. This property is sort of like the division algorithm, but not quite as strong. In the event that b is not a multiple of a, we do not insist that some multiple of a can be subtracted from b to produce an element of small valuation, but only that some linear combination of a and b has sufficiently small valuation. A valuation with this feature is called a *Dedekind-Hasse norm*, and if a domain is such that there exists a Dedekind-Hasse norm, then you can show the following.

EXERCISE 9.11.8 Suppose D is a domain with a valuation d and with the property that, for all nonzero a and b, either $b \in (a)$, or there exist domain elements m and n such that $d(ma + nb) < d(a)$. Then D is a PID.[11]

The trick to showing that $\mathbb{Z}[\alpha]$ is a PID is to show that $\mathbb{Z}[\alpha]$ entertains a Dedekind-Hasse norm. A valuation that works is $d(a + b\alpha) = a^2 + ab + 5b^2$, for which we will not provide any details.

9.12 Polynomials over a Field

An important example of a Euclidean domain is $\mathbb{Q}[t]$, and we want to address this now. Every nonzero polynomial has a nonnegative degree, which is precisely the valuation we will use. Since $\deg(fg) = \deg f + \deg g$, and since $\deg g \geq 0$, we have that $\deg f \leq \deg(fg)$. Showing $\mathbb{Q}[t]$ is Euclidean then boils down to proving a sort of division algorithm. So here is the theorem we need, with uniqueness of q and r to boot. The technique we use to get the proof off the ground should look surprisingly familiar. You will provide a few of the details in an exercise to follow.

Theorem 9.12.1 Suppose f and g are nonzero polynomials in $\mathbb{Q}[t]$. Then there exist unique polynomials q and r such that $g = fq + r$ and either $r = 0$ or $\deg r < \deg f$.

[11] Show that an arbitrary ideal I is principal by letting $a \in I$ be such that $d(a)$ is minimum. You can then show that any $b \in I$ must be a multiple of a.

Proof. Let f and g be nonzero polynomials, and define

$$S = \{g - fq : q \in \mathbb{Q}[t]\} \tag{9.57}$$

If the zero polynomial is in S, then we have $g = fq$. Otherwise we may let r be any polynomial in S whose degree is minimal. Then $r = g - fq$ for some polynomial q, so that $g = fq + r$. To show $\deg r < \deg f$, suppose $\deg r \geq \deg f$. Then we may write

$$f = a_m t^m + \cdots a_0 \quad \text{and} \quad r = b_n t^n + \cdots b_0 \tag{9.58}$$

where $m \leq n$, and neither a_m nor b_n is zero. By Exercise 9.12.2, we may create a nonzero element of S whose degree is strictly less than the degree of r, which is a contradiction. Also by Exercise 9.12.2, q and r are unique. □

Exercise 9.12.2 Finish the proof of Theorem 9.12.1 by showing the following.

(a) The polynomial $r_1 = g - fq - (b_n/a_m)t^{n-m}f$ is an element of S such that $\deg r_1 < \deg r$.

(b) If $g = fq_1 + r_1$ and $g = fq_2 + r_2$ where $r_1 = 0$ or $\deg r_1 < \deg f$ and where $r_2 = 0$ or $\deg r_2 < \deg f$, then $q_1 = q_2$ and $r_1 = r_2$.[12]

Having shown that $\mathbb{Q}[t]$ is Euclidean opens a floodgate of interesting facts about this very important ring. First, unlike $\mathbb{Z}[t]$, $\mathbb{Q}[t]$ is a PID, so every nontrivial ideal has a nonzero generator, and its degree is minimal among all elements of the ideal. Furthermore, from our comments before the statement of Exercise 9.11.6, *any* polynomial in the ideal whose degree is minimal will generate the ideal. If we write $f = a_n t^n + a_{n-1}t^{n-1} + \cdots + a_0$, then $a_n \neq 0$ and we can create an associate polynomial $m = a_n^{-1} f = t^n + (a_{n-1}/a_n)t^{n-1} + \cdots + a_0/a_n$ whose leading coefficient is one. Since m and f have the same degree, they generate the same ideal. A polynomial whose leading coefficient is the unity element of the ring of coefficients is called *monic*, and we have therefore argued the following.

Theorem 9.12.3 Every ideal in $\mathbb{Q}[t]$ has a monic generator.

Since $\mathbb{Q}[t]$ is Euclidean, it is also a UFD, so that every nonzero polynomial factors into irreducible polynomials. Furthermore, this factorization is unique, at least to a point. Any two factorizations that appear to be different can be seen upon inspection to have components that pair up as associates. Now associates are unit multiples of each other, and from Exercise 9.8.24(b), the units in $\mathbb{Q}[t]$ are

[12] What can you say about $\deg(r_2 - r_1)$? What does Exercise 9.8.22 allow you to conclude?

the polynomials of degree zero. Thus two polynomials are associates if and only if one is a nonzero constant multiple of the other.

Example 9.12.4 Two factorizations of t^3+3t^2+2t are $t(t+2)(t+1)$ and $(4t)(\frac{1}{3}t+\frac{2}{3})(\frac{3}{4}t+\frac{3}{4})$. Associate pairs are t and $4t$, $t+2$ and $\frac{1}{3}t+\frac{2}{3}$, and $t+1$ and $\frac{3}{4}t+\frac{3}{4}$. ■

We have said a lot about irreducible polynomials in $\mathbb{Q}[t]$, but we have not developed criteria by which we can determine whether a polynomial is reducible or irreducible. By Exercise 9.8.22, if $\deg f=1$ and $f=gh$, then either g or h has degree zero and is therefore a unit. Thus polynomials of degree one are irreducible.

Another simple criterion for irreducibility involves evaluating the polynomial at some rational number. In Section 9.4, we said we were not really interested in polynomials as functions where we would plug in values for t, but more as a string of symbols. Well, that statement is now false. As a ring in and of itself, a polynomial ring is exactly as we described it in Section 9.4, and t is indeed just a formal (and foreign) symbol whose presence is used to define the addition and multiplication of two elements. But for a given polynomial f there are some pretty good reasons we might want to choose a specific ring element α and calculate the value of $f(\alpha)$. The ring element $f(\alpha)$ can make an important statement about the polynomial f and its role in the polynomial ring.

You spent time in high school algebra trying to factor polynomials, and one way you might have stumbled onto a linear factor of f was by discovering some number a such that $f(a)=0$. For example, if $f=t^3-t^2+2t-2$, you might have noticed that $f(1)=0$ by adding the coefficients, so you could write $f=(t-1)g$, which by division became $f=(t-1)(t^2+2)$. The reason this worked was the following exercise.

EXERCISE 9.12.5 Suppose f is a polynomial in $\mathbb{Q}[t]$ and a is a rational number. Then $(t-a)\mid f$ if and only if $f(a)=0$.

Exercise 9.12.5 makes the following almost immediate.

Theorem 9.12.6 Suppose f is a polynomial in $\mathbb{Q}[t]$ whose degree is either 2 or 3. Then f is reducible if and only if there exists a rational number a such that $f(a)=0$.

Proof. Let f be a polynomial in $\mathbb{Q}[t]$ of degree 2 or 3.

($\Rightarrow$) Suppose f is reducible. Then $f=gh$, where the degrees of g and h are nonzero, but strictly less than $\deg f$. Without loss of generality, we may assume $\deg g=1$, so that $g=at+b$ for rational numbers a and b, where a is not zero. Thus $f(-b/a)=g(-b/a)h(-b/a)=0$.

($\Leftarrow$) Suppose there exists rational a such that $f(a)=0$. By Exercise 9.12.5, $(t-a)\mid f$, so that f is reducible. □

Whether a polynomial with rational coefficients is reducible is not an easy question to answer in general. Several criteria can help answer the question for certain special polynomials, and you will see them in your upper level algebra class.

If there is one important idea that you should have clued into by this point in the mathematical game, it is that certain properties of mathematical structures are the logical basis for other properties and that these properties can sometimes be isolated and translated over to other structures that might be very different. For example, we proved in Theorem 2.1.11 that $a \cdot 0 = 0$ for all integers a. Then we noted in Section 9.1 that $a \cdot 0 = 0 \cdot a = 0$ in any ring by an argument identical to that for the integers. A few basic ring properties come together to make this true (basic properties of addition, including additive cancellation, and the distributive property), even without commutativity of multiplication. So any mathematical structure in which we have additive cancellation and the distributive property will be a structure where $a \cdot 0 = 0 \cdot a = 0$.

The fact that $\mathbb{Q}[t]$ is a Euclidean domain calls on the fact that the rational numbers are a field, but it does not require that they be any particular kind of field. If we replace the rationals with an arbitrary field K in Theorems 9.12.1–9.12.6, exactly the same proofs work. This can be particularly interesting if we use the field $\mathbb{Z}_p$ for prime number p. Here are restatements of Theorems 9.12.1–9.12.6 for an arbitrary field and some examples to illustrate their broader application.

Theorem 9.12.7 Let K be a field. Then $K[t]$ with valuation $\deg f$ is a Euclidean domain.

Exercise 9.12.8

(a) Let $f_1 = 3t^2 + t + 2$ and $g_1 = 3t^4 + 3t^3 + t^2 + 2t + 4$ be polynomials in $\mathbb{Z}_5[t]$. Show that there exist q and r in $\mathbb{Z}_5[t]$ such that $g_1 = f_1 q + r$ and $\deg r < \deg f_1$.[13]

(b) Let $f_2 = 2t + 2$ and $g_2 = t^2$ be polynomials in $\mathbb{Z}_6[t]$. Show that it is impossible to write $g_2 = f_2 q + r$ for any q and r in $\mathbb{Z}_6[t]$ where either $r = 0$ or $\deg r < \deg f_2$.[14]

Theorem 9.12.9 Let K be a field, $f \in K[t]$ and $a \in K$. Then $(t - a) \mid f$ if and only if $f(a) = 0$.

Theorem 9.12.10 Let K be a field, $f \in K[t]$ and suppose the degree of f is either 2 or 3. Then f is reducible if and only if there exists $a \in K$ such that $f(a) = 0$.

[13] Use polynomial division.

[14] The coefficient of t^2 in $f_2 q$ must be one. Show this is impossible.

EXERCISE 9.12.11 Apply Theorem 9.12.10 to the following polynomials in $\mathbb{Z}_3[t]$ by either finding a proper factorization or explaining why they are irreducible.

(a) $f_1 = t^2 + t + 1$

(b) $f_2 = t^2 + t + 2$

(c) $f_3 = t^3 + t^2 + 2$

(d) $f_4 = t^3 + t + 2$

9.13 Polynomials over the Integers

We have waited to show that $\mathbb{Z}[t]$ is a UFD, and now is the time to tackle the question. We could have shown the existence of a factorization of a polynomial in $\mathbb{Z}[t]$ into irreducibles earlier, but uniqueness of this factorization up to order and association of the factors requires us to view elements of $\mathbb{Z}[t]$ as elements of $\mathbb{Q}[t]$, where factorizations are unique up to association. The reason we have some work to do to show uniqueness is that reducibility and association in $\mathbb{Q}[t]$ are different from reducibility and association in $\mathbb{Z}[t]$. Polynomials such as $2t + 6$ and $10t + 15$ are irreducible and associates in $\mathbb{Q}[t]$, but not in $\mathbb{Z}[t]$, because the constant polynomials 2 and 5 are units in $\mathbb{Q}[t]$ but not in $\mathbb{Z}[t]$. First the easy part.

EXERCISE 9.13.1 Every nonzero, non-unit polynomial in $\mathbb{Z}[t]$ has a factorization into irreducible polynomials in $\mathbb{Z}[t]$.[15]

To make our way to uniqueness up to order and association, we need the following lemma. You will provide the climactic detail as an exercise. Remember that the term *primitive* applies only to polynomials of degree at least one.

Theorem 9.13.2 (Gauss's Lemma). In $\mathbb{Z}[t]$, the product of two primitive polynomials is primitive.

Proof. Suppose f and g are primitive polynomials in $\mathbb{Z}[t]$. We show that fg is primitive by supposing p is any prime number, and then showing there is some coefficient in fg that is not divisible by p.

Suppose p is a prime number, and write

$$f = a_m t^m + \cdots + a_0 \quad \text{and} \quad g = b_n t^n + \cdots + b_0 \tag{9.59}$$

[15] Use strong induction on the degree of f and mimic the proof of Theorem 3.5.19. Exercise 9.10.1 takes care of the case $\deg(f) = 0$.

Since f and g are primitive, there are coefficients in both f and g that are not divisible by p. Let a_j and b_k be, respectively, the coefficients of the lowest powers of t in f and g that are not divisible by p; that is, p divides all of $a_0, a_1, \ldots, a_{j-1}, b_0, b_1, \ldots, b_{k-1}$, but not a_j or b_k. Then by Exercise 9.13.3, p does not divide the coefficient of t^{j+k} in fg. Since this is true for all prime numbers, fg is primitive. □

EXERCISE 9.13.3 Finish the proof of Theorem 9.13.2 by showing p does not divide the coefficient of t^{j+k} in fg.[16]

The next theorem says simply that if a polynomial has integer coefficients, and it can be factored into polynomials of degree at least one by viewing it as an element of $\mathbb{Q}[t]$ and resorting to rational coefficients in the factors, then you can adjust these factor coefficients and make them integers. We will prove this in a somewhat conversational way to keep the notation from getting too sloppy. Notice the point at which we apply Exercise 9.13.3.

Theorem 9.13.4 Suppose f is a polynomial with integer coefficients. If f is reducible in $\mathbb{Q}[t]$, then it is reducible in $\mathbb{Z}[t]$.

Proof. Suppose f has integer coefficients, and suppose $f = f_1 f_2$, where f_1 and f_2 are polynomials with rational coefficients, and both f_1 and f_2 have degree at least 1. Let d_1 and d_2 be, respectively, the product of all the denominators of all the coefficients of f_1 and f_2; then let $g_1 = d_1 f_1$ and $g_2 = d_2 f_2$, so that g_1 and g_2 have integer coefficients. Next, let c_1 and c_2 be the content of g_1 and g_2, respectively, and factor these out to write $g_1 = c_1 h_1$ and $g_2 = c_2 h_2$, so that h_1 and h_2 are primitive. Also, if we let c be the content of f, we may write $f = cg$, where g is primitive. Thus we have

$$(cd_1d_2)g = (d_1d_2)f = d_1 f_1 d_2 f_2 = g_1 g_2 = (c_1c_2)h_1h_2 \tag{9.60}$$

where g, h_1, and h_2 are all primitive polynomials. By Exercise 9.13.3, h_1h_2 is also primitive, so the content of the left-hand side of Eq. (9.60) is cd_1d_2, and the content of the right-hand side is c_1c_2. Since cd_1d_2 and c_1c_2 are both positive integers, $cd_1d_2 = c_1c_2$, and $g = h_1h_2$. Thus $f = cg = ch_1h_2$, and we have a proper factorization of f with integer coefficients. □

Now the result we have all been waiting for.

Theorem 9.13.5 The factorization in Theorem 9.13.4 is unique up to order and association of the factors.

[16] The coefficient of t^{j+k} is $\sum_{i=0}^{j+k} a_i b_{j+k-i}$. Look separately at the terms $0 \le i \le j-1$, $i = j$, and $j+1 \le i \le j+k$.

Proof. Suppose $f \in \mathbb{Z}[t]$ can be written as

$$f = p_1 p_2 \dots p_m = q_1 q_2 \dots q_n \tag{9.61}$$

where all p_k and q_k are irreducible polynomials with integer coefficients. If $n = 1$, then $m = 1$ and $p_1 = q_1$, so that the result is trivially true.

So suppose $n \geq 2$ and that the result is true for all polynomials of degree less than n. Note that all p_k and q_k are either prime constant polynomials or primitive. Since $p_m \mid q_1 \dots q_n$, then viewing p_m and all q_k as elements of $\mathbb{Q}[t]$, irreducible in $\mathbb{Q}[t]$ by Theorem 9.13.4, we have that $p_m \mid q_k$ (in $\mathbb{Q}[t]$) for some k by Theorem 9.10.12. Reordering the q_k, we may assume $p_m \mid q_n$, and since q_n is irreducible in $\mathbb{Q}[t]$, we may write $(a/b)p_m = q_n$ for some rational number a/b. Thus $ap_m = bq_n$, and since p_m and q_n are primitive, $a = \pm b$. Therefore $q_n = \pm p_m$, so that p_m and q_n are associates in $\mathbb{Z}[t]$. Substituting $\pm p_m$ for q_n in Eq. (9.61) and canceling p_m, we have

$$p_1 \dots p_{m-1} = \pm q_1 \dots q_{n-1} \tag{9.62}$$

By the inductive assumption $m = n$, and the remaining p_k and q_k may be reordered and paired as associates, so that the factorization of f into irreducible polynomials with integer coefficients is unique up to order and association of the factors. □

9.14 Ring Morphisms

The theory of ring morphisms should seem a breezy topic after having studied group morphisms in Section 8.6. The only difference in the definition in the context of rings is that there are two binary operations to preserve. As we are doing increasingly often, we will be fairly relaxed about notation for elements and operations in the two rings, unless we need it to avoid confusion.

Definition 9.14.1 Suppose R and S are rings and $\phi : R \to S$ is a function with the properties that $\phi(x + y) = \phi(x) + \phi(y)$ and $\phi(xy) = \phi(x)\phi(y)$ for all $x, y \in R$. Then ϕ is called a *homomorphism* (or *morphism*) from R to S. The terms *monomorphism*, *epimorphism*, *isomorphism*, and *automorphism* are defined in a way analogous to that for groups, and if there exists an isomorphism from R to S, we write $R \cong S$.

Here are some examples. Let's start with a trivial one.

Example 9.14.2 If R and S are rings, the mapping defined by $\phi(x) = 0$ for all $x \in R$ is called the *trivial* morphism. ■

Example 9.14.3 The identity mapping is an automorphism of any ring. ■

Example 9.14.4 Define $\phi : \mathbb{Z} \to \mathbb{Z}_n$ by $\phi(x) = (n) + x$, the equivalence class of $x \mod n$. If we think of $\mathbb{Z}_n$ as the set of integers $\{0, 1, \ldots, n-1\}$ with clock arithmetic, ϕ maps x to its remainder upon division by n, and ϕ is an epimorphism. For

$$\phi(x + y) = (n) + [x + y] = [(n) + x] + [(n) + y] = \phi(x) + \phi(y) \tag{9.63}$$

$$\phi(xy) = (n) + xy = [(n) + x][(n) + y] = \phi(x)\phi(y) \tag{9.64}$$

Furthermore, ϕ is onto. For if $(n) + k$ is any coset in $\mathbb{Z}_n$, then $\phi(k) = (n) + k$. ■

Example 9.14.5 For $\mathbb{Z}[\sqrt[n]{x}]$ in the form of Eq. (9.33), define $\phi : \mathbb{Z} \to \mathbb{Z}[\sqrt[n]{x}]$ by

$$\phi(k) = k + 0\sqrt[n]{x} + \cdots + 0\sqrt[n]{x^{n-1}} \tag{9.65}$$

Then ϕ is a monomorphism. ■

Example 9.14.5 illustrates a subtle distinction between two ideas. After we constructed $\mathbb{Z}[\sqrt[3]{2}]$ in Example 9.4.1, we said (page 303) that the integers are a subring of $\mathbb{Z}[\sqrt[3]{2}]$. If we had been a bit more rigorous in our construction of $\mathbb{Z}[\sqrt[3]{2}]$, we would have built it up, not as a set of expressions of the form $a + b\sqrt[3]{2} + c\sqrt[3]{4}$, but as a completely formal set of ordered triples

$$\mathbb{Z}[\sqrt[3]{2}] = \{(a, b, c) : a, b, c \in \mathbb{Z}\} \tag{9.66}$$

where addition and multiplication are defined to coincide with the definitions in Example 9.4.1. Specifically, defining addition and multiplication by

$$\begin{aligned} (a, b, c) + (d, e, f) &= (a + d, b + e, c + f) \\ (a, b, c)\cdot(d, e, f) &= (ad + 2bf + 2ce, ae + bd + 2cf, af + be + cd) \end{aligned} \tag{9.67}$$

incorporates the behavior of $\sqrt[3]{2}$ into the operations, even though they are mere manipulations of ordered triples with no apparent presence of $\sqrt[3]{2}$. With these definitions, to say that the integers are a subring of $\mathbb{Z}[\sqrt[3]{2}]$ is technically not true, for $\mathbb{Z}$ is not a subset of $\mathbb{Z}[\sqrt[3]{2}]$ as subrings must be. Whereas the integers are a set of numbers, $\mathbb{Z}[\sqrt[3]{2}]$ is a set of ordered triples of numbers. However, it does not mean that the link between them is somehow illusory. Defining $\psi : \mathbb{Z} \to \mathbb{Z}[\sqrt[3]{2}]$ by $\psi(n) = (n, 0, 0)$, we have an exact parallel to ϕ in Eq. (9.65). Whether you think of $\mathbb{Z}[\sqrt[3]{2}]$ as we originally defined it in Example 9.4.1 or as in Eqs. (9.66) and (9.67), we say that the integers have been *embedded monomorphically* in $\mathbb{Z}[\sqrt[3]{2}]$. Imagine the integers and $\mathbb{Z}[\sqrt[3]{2}]$ as separate, where $\mathbb{Z}$ is a set of elements of the form n and $\mathbb{Z}[\sqrt[3]{2}]$ consists of elements of the form (a, b, c). Then the range of ϕ is the set of all elements of $\mathbb{Z}[\sqrt[3]{2}]$ of the form $(a, 0, 0)$ and is isomorphic to the integers, and hence structurally the same.

Example 9.14.5 illustrates a slight breach of rigor we committed in defining ring extensions in Section 9.4, so let's clear that up. It illustrates a technicality when we say that $\mathbb{Z}[\sqrt[3]{2}]$ is an extension of the integers. If R and S are rings where R is a subset of S, then saying S is an extension of R has the same meaning as saying R is a subring of S. But if you start with R and want to extend it to some S by adjoining an element or elements, the standard, more rigorous way is to build S from scratch and then monomorphically embed R in S. Then when we say that S is an extension of R, we mean that S is the range of a monomorphism whose domain is R.

EXERCISE 9.14.6 Let R be a commutative ring, and fix some ring element α. Define $\phi_\alpha : R[t] \to R$ by $\phi(f) = f(\alpha)$. Clearly, ϕ_α is a function (quick mental exercise). Show that ϕ is an epimorphism. This particularly important morphism is called the *evaluation at* α morphism. It maps every polynomial in $R[t]$ to its value at α.

EXERCISE 9.14.7 Let R be a commutative ring with unity. Describe the evaluation morphisms ϕ_0 and ϕ_e.

EXERCISE 9.14.8 On the Gaussian integers, define $\phi(a + bi) = a - bi$. Show that ϕ is an automorphism. This is called the *conjugation* morphism, for it sends a Gaussian integer to its *complex conjugate*.

EXERCISE 9.14.9 Let R be a ring with unity e, and define $\phi : \mathbb{Z} \to R$ by $\phi(n) = ne$. Show that ϕ is a morphism.

EXERCISE 9.14.10 Define $\phi : \mathbb{Z}_{2\times 2} \to \mathbb{Z}$ by $\phi\left(\begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}\right) = a_{11}$. Is ϕ a morphism? Prove or disprove.

EXERCISE 9.14.11 Define $\phi : \mathbb{Z} \to \mathbb{Z}_{2\times 2}$ by $\phi(n) = \begin{bmatrix} n & 0 \\ 0 & 0 \end{bmatrix}$. Is ϕ a morphism? Prove or disprove.

9.14.1 Properties of Ring Morphisms

Since a ring morphism preserves addition between the two rings as abelian additive groups, our results from Section 8.6 apply to addition. Thus we gain the following for free from Theorem 8.6.8 and Exercise 8.6.10.

Theorem 9.14.12 Suppose $\phi : R \to S$ is a ring morphism. Then

1. $\phi(0) = 0$.
2. $\phi(-r) = -\phi(r)$ for all $r \in R$.
3. $\phi(nr) = n\phi(r)$ for all $r \in R$ and all integers n.

To have analogous results for multiplication, we have to make some modifications.

EXERCISE 9.14.13 If R is a ring with unity e_R and if $\phi : R \to S$ is a ring morphism such that $\phi(e_R) = 0$, then ϕ is the trivial morphism.

The contrapositive of Exercise 9.14.13 reveals that if $\phi : R \to S$ is a nontrivial morphism, then the unity of R does not map to the zero element of S. If S has unity e_S, then comparable to Theorem 8.6.8, you might be tempted to think that a nontrivial morphism $\phi : R \to S$ would have to satisfy $\phi(e_R) = e_S$. To prove such a property and the exponent rules related to it, multiplicative cancellation is necessary in S. Then a proof of the next result becomes possible. Since the proof of the exponent rules in parts 2 and 4 would be identical to the proof of Exercise 8.6.10, you only need to prove parts 1 and 3.

Theorem 9.14.14 Suppose R is a ring with unity e_R, S is an integral domain with unity e_S, and $\phi : R \to S$ is a nontrivial ring morphism.

1. $\phi(e_R) = e_S$.
2. For all $r \in R$ and nonnegative integers n, $\phi(r^n) = [\phi(r)]^n$.
3. If $u \in R$ is a unit, then so is $\phi(u)$, and $\phi(u^{-1}) = [\phi(u)]^{-1}$.
4. If $u \in R$ is a unit, then $\phi(u^{-n}) = [\phi(u)]^{-n}$ for all positive integers n.

EXERCISE 9.14.15 Prove parts (1) and (3) of Theorem 9.14.14.

Theorems 9.14.12 and 9.14.14 reveal that the only nontrivial morphism from the integers to any other integral domain is defined by $\phi(n) = ne$. For suppose $\psi : \mathbb{Z} \to D$ is a nontrivial morphism and e is the unity in D. Since 1 is the unity in the integers, it must be that $\psi(1) = e$. Furthermore, for any integer n, $\psi(n) = \psi(n \cdot 1) = n\psi(1) = ne$. Thus the only nontrivial morphism from the integers to the integers is the identity. You can use other parts of Theorems 9.14.12 and 9.14.14 to show the following.

EXERCISE 9.14.16 The only automorphism of the rationals is the identity.

If S is not a domain, Theorem 9.14.14 might not apply. For example, if ϕ is the morphism from Exercise 9.14.11, $\phi(1) = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$, which is not the identity matrix. Furthermore, though 1 is a unit in the integers, $\phi(1)$ is not a unit in $\mathbb{Z}_{2\times 2}$.

A lot of our work with group morphisms involved their relationship to normal subgroups. The interesting relationships between ring morphisms and substructures involve ideals. The next theorem is a parallel to Exercise 8.6.14. When you prove the left-sided case as an exercise, parts of your proof will follow immediately from your work in group theory and not require new arguments. For example, in your proof of part (1), you may simply point out that since R is an additive group with respect to addition, $\phi(R)$ is an additive subgroup of S by Exercise 8.6.14. That takes care of three of the requirements in showing that $\phi(R)$ is a subring of S.

Theorem 9.14.17 Suppose $\phi : R \to S$ is a ring morphism. Then

1. $\phi(R)$ is a subring of S.
2. If I is a left (right) ideal of R, then $\phi(I)$ is a left (right) ideal of $\phi(R)$.
3. If I is a left (right) ideal of S, then $\phi^{-1}(I)$ is a left (right) ideal of R.

Exercise 9.14.18 Prove the left-sided cases of Theorem 9.14.17.

Similar to groups, the kernel of a ring morphism is defined as

$$\text{Ker}(\phi) = \{r \in R : \phi(r) = 0\} \tag{9.68}$$

Right away, we have the following. Save yourself some work by applying Exercise 8.6.16.

Exercise 9.14.19 If $\phi : R \to S$ is a ring morphism, then $\text{Ker}(\phi)$ is a two-sided ideal of R.

The following needs no additional proof, for it follows from Exercise 8.6.17 applied to R as an additive group.

Theorem 9.14.20 If $\phi : R \to S$ is a ring morphism, then $\text{Ker}(\phi) = \{0\}$ if and only if ϕ is one-to-one.

If R is any ring with unity element e, the morphism from the integers defined by $\phi(n) = ne$ can look either of two ways, depending on the characteristic of R. If R has characteristic zero, then by definition ne is never zero for nonzero integer n. The following should therefore be quick.

Exercise 9.14.21 If R is a ring with unity e and characteristic zero, then $\phi : \mathbb{Z} \to R$ defined by $\phi(n) = ne$ is one-to-one.

Exercises 9.14.9 and 9.14.21 say that the integers can be monomorphically embedded in a ring R with unity and characteristic zero. We loosely say that R *contains* the integers, meaning it contains a subring generated by its unity that is isomorphic to the integers. For example, in $\mathbb{R}_{2\times 2}$, the image of the integers under ϕ is

$$\left\{\ldots, \begin{bmatrix} -2 & 0 \\ 0 & -2 \end{bmatrix}, \begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix}, \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}, \ldots\right\} \tag{9.69}$$

If R has nonzero characteristic, then there is a smallest positive integer n for which $ne = 0$. It might seem pretty clear that R in this case will contain a subring that looks like $\mathbb{Z}_n$ instead of the integers. Before we can make an argument for this, however, we need to create the notion of a *quotient ring*, which we will do in the next section.

9.15 Quotient Rings

With a group and a normal subgroup we can build a quotient group. In an analogous way, given a ring and an ideal, we can build a quotient ring. Part of the work in building a quotient ring is exactly like building a quotient group, so the fact that a quotient ring has properties R1–R10 has already been done in part. As with groups, building the quotient structure begins by defining a form of equivalence.

Theorem 9.15.1 Let R be a ring and I an ideal of R. For ring elements a and b, define $a \equiv_I b$ if $a - b \in I$. Then $\equiv_I$ is an equivalence relation on R.

Since R is an abelian group with respect to its addition operation, I is a normal additive subgroup of R. Therefore, Theorem 9.15.1 is merely a restatement of Exercise 8.3.2 in its additive form and needs no additional proof. So we are ready to define the set R/I with its addition and multiplication operations, and show that it has all properties R1–R10. There are no surprises in the definitions of the binary operations on R/I. But considering the need for H to be a normal subgroup of G in showing the binary operation on G/H is well defined, little bells should be going off in your head as you consider the burden of showing that addition and multiplication on R/I are well defined. It should be no surprise that you will exploit the fact that I is an ideal of R in at least part of this demonstration. What you might not see at first is what kind of ideal I needs to be and where you will call on the fact that I is this kind of ideal. Since addition in R is commutative, I is normal as an additive subgroup of R. Thus addition is well defined by our work in Chapter 8. It's a different story for multiplication, though. If multiplication is not commutative, then a left ideal might not be a right ideal, and vice versa. So we might wonder whether I needs merely to be either a left ideal or a right ideal, or perhaps both. It turns out that I needs to be a two-sided ideal, and you will see why when you provide the missing piece of the proof of the following.

Theorem 9.15.2 Let R be a ring and I a two-sided ideal of R. For $a \in R$ write $I + a = [a]$, where $[a]$ is the equivalence class of a modulo I. Define addition $\oplus$ and multiplication $\otimes$ on $R/I = \{I + a : a \in R\}$ by

$$(I + a) \oplus (I + b) = I + (a + b) \quad \text{and} \quad (I + a) \otimes (I + b) = I + (ab) \tag{9.70}$$

Then R/I is a ring under the operations $\oplus$ and $\otimes$.

We can talk our way through almost all properties R1–R10, so that your work in proving Theorem 9.15.2 will be minimal. Since R is an abelian additive group and I is a normal additive subgroup, R/I has properties R1–R6. Properties R8–R10 are immediate from the definitions of $\oplus$ and $\otimes$. The only property that takes any real work is R7, showing that $\otimes$ is well defined, and this is what you will show in the next exercise. True to form, you suppose $I + a = I + b$ and $I + c = I + d$ and use this to show that $I + ac = I + bd$. That is, supposing $a - b$ and $c - d$ are in I should somehow allow you to show that $ac - bd$ is in I also.

Exercise 9.15.3 Finish the proof of Theorem 9.15.2 by showing that multiplication is well defined.

With R/I defined and shown to be a ring, we should be able to bypass all the chatty exposition as in Chapter 8 and jump right to results analogous to Theorems 8.6.19 and 8.6.21.

Exercise 9.15.4

(a) State a result analogous to Theorem 8.6.19 for a ring and the quotient ring created by "modding" out an ideal.

(b) Explain how all except one component of the proof of this theorem follows from Theorem 8.6.19.

(c) Provide the missing component to complete the proof of your theorem.

Exercise 9.15.5 Suppose R and S are rings and $\phi : R \to S$ is an epimorphism. Then $S \cong R/\operatorname{Ker}(\phi)$.

Let's return to the morphism in Exercise 9.14.9 defined by $\phi(n) = ne$. If a ring has nonzero characteristic, then there is a smallest positive integer n for which $ne = 0$. By Exercise 9.14.19, $\operatorname{Ker}(\phi)$ is an ideal in the integers, which is a PID. Thus $\operatorname{Ker}(\phi) = (k)$ for some integer k. Since $\phi(k) = 0$, and since n is the smallest positive integer for which $ne = 0$, it must be that $k \geq n$. But since $n \in \operatorname{Ker}(\phi)$, it must be that n is a multiple of k, so that $k \leq n$. Thus $k = n$. By Exercise 9.15.5, the range of ϕ is isomorphic to $\mathbb{Z}/(n)$. That is, R contains a subring isomorphic to $\mathbb{Z}_n$.

If we think of a ring with unity e as an abelian additive group and let S be the subgroup generated by e, the identity element of the *other* operation, then the form of S can be determined by looking at the additive form of Eq. (8.22).

$$S = \{ne : n \in \mathbb{Z}\} \tag{9.71}$$

which is precisely the range of ϕ. Thus as far as addition is concerned, S is the smallest additive subgroup of R that contains e. If we could show that S is closed under multiplication, we would have that S is the smallest subring of R that contains e. But closure is immediate by Exercise 9.3.20. In bits and pieces, we have proved the following.

Theorem 9.15.6 Suppose R is a ring with unity. If R has characteristic zero, then it contains a subring isomorphic to the integers. If R has positive characteristic n, then it contains a subring isomorphic to $\mathbb{Z}_n$. In either case, such is the smallest subring of R that contains e.

Because the polynomial ring over a field is Euclidean, it makes for some particularly important quotient rings. Let's spend some time studying the quotient ring created by "modding" out the ideal generated by $f = 2t^3 - t + 5$ from $\mathbb{Q}[t]$. The ties back to the integers and the integers modulo n are uncanny. To be concrete, we will let $n = 6$ and draw parallels between the relationship of the integers to the integers modulo 6 and the relationship of $\mathbb{Q}[t]$ to $\mathbb{Q}[t]/(f)$.

Since the integers are a Euclidean domain, every integer can be written as $6q + r$ for integers q and r and where $0 \leq r \leq 5$. Consequently, every integer is equivalent mod 6 to some element of $\{0, 1, 2, 3, 4, 5\}$, so that every element of $\mathbb{Z}_6$ is a coset that can be addressed by a unique representative element in $\{0, 1, 2, 3, 4, 5\}$. To perform addition and multiplication in $\mathbb{Z}_6$, a purist would write something like

$$[(6) + 4] + [(6) + 3] = (6) + 7 = (6) + 1 \tag{9.72}$$

or

$$[(6) + 5] \times [(6) + 4] = (6) + 20 = (6) + 2 \tag{9.73}$$

where the first step in these two calculations is an application of the definitions of addition and multiplication in $\mathbb{Z}_6$ from Theorem 9.15.2 and the second step is an application of equivalence mod 6 to simplify the calculation to a standard form with representative element from $\{0, 1, 2, 3, 4, 5\}$. As long as we realize that this is what we're doing, we can write these calculations as

$$4 + 3 =_6 7 =_6 1 \quad \text{and} \quad 5 \times 4 =_6 20 =_6 2 \tag{9.74}$$

This form has the imagery of performing addition and multiplication in the integers, with the stipulation that any sum or product that gets kicked out of bounds, that is, out of $\{0, 1, 2, 3, 4, 5\}$, is hauled back into $\{0, 1, 2, 3, 4, 5\}$ by subtracting a multiple of 6. All this helps us see how we may view elements of $\mathbb{Z}_6$ and how they combine by addition and multiplication to produce other elements of $\mathbb{Z}_6$.

For our specific polynomial $f = 2t^3 - t + 5$, the way to visualize elements of $\mathbb{Q}[t]/(f)$ in terms of polynomials in $\mathbb{Q}[t]$ is strikingly similar. Since $\mathbb{Q}[t]$ is a Euclidean domain, any polynomial can be written uniquely as $fq + r$ for polynomials q and r, where either $r = 0$ or $\deg r < \deg f$. So if we consider any coset $(f) + g$ in the quotient ring, there is a unique polynomial r such that $(f) + r = (f) + g$ and either $r = 0$ or $\deg r < \deg f$. Thus elements of $\mathbb{Q}[t]/(f)$ may always be addressed by a representative polynomial r where either $r = 0$ or $\deg r < \deg f$. Since we are using $f = 2t^3 - t + 5$, then we may write

$$\mathbb{Q}[t]/(f) = \{(f) + at^2 + bt + c : a, b, c \in \mathbb{Q}\} \tag{9.75}$$

and know that every element of the quotient ring can be written as some coset $(f) + at^2 + bt + c$. Furthermore, if $(f) + r_1 = (f) + r_2$, then $r_2 \equiv_{(f)} r_1$, so that $r_2 - r_1$ is a multiple of f. Now nonzero multiples of f cannot have a degree less than the degree of f, but it is impossible that $\deg(r_2 - r_1) \geq \deg f$. Thus $r_1 = r_2$, and we have that different polynomials of the form $at^2 + bt + c$ will always generate different cosets of (f).

In the same way that we view elements of $\mathbb{Z}_6$ simply as $\{0, 1, 2, 3, 4, 5\}$, we can view elements of $\mathbb{Q}[t]/(f)$ simply as polynomials of the form $at^2 + bt + c$. But we must consider how they add and multiply. Instead of using coset notation and writing $[(f) + a_1t^2 + b_1t + c_1] + [(f) + a_2t^2 + b_2t + c_2]$, we can just write

$$[a_1t^2 + b_1t + c_1] + [a_2t^2 + b_2t + c_2] =_{(f)} (a_1 + a_2)t^2 + (b_1 + b_2)t + (c_1 + c_2) \tag{9.76}$$

Adding two such polynomials cannot produce a sum of any larger degree, so Eq. (9.76) is all that needs to be said about addition in the quotient ring. However, for multiplication, let's illustrate with a concrete example, where the details of polynomial division have been omitted.

$$\begin{aligned}(4t^2 + 2)(3t^2 - 2t + 8) &=_{(f)} 12t^4 - 8t^3 + 38t^2 - 4t + 16 \\ &=_{(f)} (6t - 4)f + 44t^2 - 38t + 36 \\ &=_{(f)} 44t^2 - 38t + 36\end{aligned} \tag{9.77}$$

If we multiply two elements of the quotient ring as if they were polynomials in $\mathbb{Q}[t]$, and we produce a product of degree at least three, we can apply the division algorithm to subtract an appropriate multiple of f from the product to produce an equivalent polynomial of the form $at^2 + bt + c$. Now it's time for you to practice this procedure.

EXERCISE 9.15.7 In $\mathbb{Q}[t]$, let $f = t^4 + 2t + 1$. Construct the form of elements of $\mathbb{Q}[t]/(f)$, and illustrate addition and multiplication.

Now for another very interesting example. Since $\mathbb{Z}_3$ is a field, $\mathbb{Z}_3[t]$ is a Euclidean domain, and we can construct the quotient ring $\mathbb{Z}_3[t]/(f)$ for $f \in \mathbb{Z}_3[t]$ in a similar way. Let's use $f = t^3 + t + 2$, construct the quotient ring, and look at addition and multiplication. By exactly the same reasoning as before, $\mathbb{Z}_3[t]/(f) = \{at^2 + bt + c : a, b, c \in \mathbb{Z}_3\}$. Notice this is a finite set. Each of a, b, and c can take on values from $\{0, 1, 2\}$, so $\mathbb{Z}_3[t]/(f)$ has 27 elements. Adding elements of $\mathbb{Z}_3[t]/(f)$ is easy:

$$(2t^2 + t + 2) + (t^2 + 2t + 2) =_{(f)} 3t^2 + 3t + 4 =_{(f)} 1 \tag{9.78}$$

Doing multiplication would look like the following if we simplify the product by way of the division algorithm.

$$\begin{aligned}(2t^2 + t + 2)(t^2 + 2t + 2) &=_{(f)} 2t^4 + 5t^3 + 8t^2 + 6t + 4 \\ &=_{(f)} 2t^4 + 2t^3 + 2t^2 + 1 \\ &=_{(f)} (2t + 2)f \\ &=_{(f)} 0\end{aligned} \tag{9.79}$$

However, there is a slick way to simplify multiplication by making substitutions. In $\mathbb{Z}_3[t]/(f)$, $f \equiv_{(f)} 0$, or $t^3 + t + 2 \equiv_{(f)} 0$. This can also be written as $t^3 \equiv_{(f)} -t - 2 \equiv_{(f)} 2t + 1$. This means that any t^3 produced in the process of multiplication can be replaced with $2t + 1$, thus bringing the degree of a product back down.

$$\begin{aligned}(2t^2 + t + 2)(t^2 + 2t + 2) &=_{(f)} 2t^4 + 5t^3 + 8t^2 + 6t + 4\\ &=_{(f)} 2t^4 + 2t^3 + 2t^2 + 1\\ &=_{(f)} (2t)t^3 + 2t^3 + 2t^2 + 1\\ &=_{(f)} (2t)(2t + 1) + 2(2t + 1) + 2t^2 + 1 \qquad (9.80)\\ &=_{(f)} 4t^2 + 2t + 4t + 2 + 2t^2 + 1\\ &=_{(f)} 6t^2 + 6t + 3\\ &=_{(f)} 0\end{aligned}$$

EXERCISE 9.15.8 In $\mathbb{Z}_3[t]$, let $f = t^4 + 2t + 1$. Construct the form of elements of $\mathbb{Z}_3[t]/(f)$, and illustrate addition and multiplication using the fact that $t^4 =_{(f)} t + 2$.

EXERCISE 9.15.9 Every element of $\mathbb{Q}[t]/(t^2 - 2)$ can be written in the form $at + b$ for rational a and b. Use a trick similar to that in Exercise 9.15.8 to simplify $(at + b)(ct + d)$.

EXERCISE 9.15.10 In Exercise 9.4.8 you showed $\mathbb{Q}[\sqrt{2}]$ is a field. Calculate and simplify $(b + a\sqrt{2})(d + c\sqrt{2})$, and compare to Exercise 9.15.9.

The last theorems and examples in this section illustrate some very interesting implications of the results we have worked so hard to develop. The next two results would not likely jump out at you as obvious, but they are elegant and not difficult to prove.

EXERCISE 9.15.11 Suppose R is a commutative ring and I is an ideal of R. Then R/I is an integral domain if and only if I is a prime ideal.

The next result could actually be stated in if-and-only-if form, but we only need one direction. If R is a commutative ring with unity e, then R/I is also a commutative ring with unity $I + e$. Also, if I is maximal, it is by definition a proper ideal of R, so that R/I has more than one element. Thus if we choose some $I + a \in R/I$ where $I + a \neq I + 0$, then a is not an element of I. Exercise 9.6.17 and the maximality of I are just what you need to prove the following.

EXERCISE 9.15.12 Suppose R is a commutative ring with unity, and I is an ideal of R. If I is maximal, then R/I is a field.

Now let's put all this together in a very elegant construction. If K is a field, then $K[t]$ is a Euclidean domain, hence a PID. If $f \in K[t]$ is irreducible (prime), then (f) is a prime ideal by Exercise 9.10.16. By Exercise 9.10.6, (f) is also maximal. Therefore, $K[t]/(f)$ is a field. We can use these facts to do the following.

Example 9.15.13 Construct a field with nine elements.

Solution Since $t^2 + 1$ has no roots in $\mathbb{Z}_3[t]$, it is irreducible. Thus $\mathbb{Z}_3[t]/(t^2 + 1) = \{at + b : a, b \in \mathbb{Z}_3\}$ is a field with nine elements. For notational simplicity, we write $at + b = (a, b)$ and illustrate multiplication in the Table 9.81. Notice $t^2 =_{(t^2+1)} 2$ and the manifestation of this in the table. Also, notice how the table reveals that every element has a multiplicative inverse.

×	(0, 0)	(0, 1)	(0, 2)	(1, 0)	(1, 1)	(1, 2)	(2, 0)	(2, 1)	(2, 2)
(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)
(0, 1)	(0, 0)	(0, 1)	(0, 2)	(1, 0)	(1, 1)	(1, 2)	(2, 0)	(2, 1)	(2, 2)
(0, 2)	(0, 0)	(0, 2)	(0, 1)	(2, 0)	(2, 2)	(2, 1)	(1, 0)	(1, 2)	(1, 1)
(1, 0)	(0, 0)	(1, 0)	(2, 0)	(0, 2)	(1, 2)	(2, 2)	(0, 1)	(1, 1)	(2, 1)
(1, 1)	(0, 0)	(1, 1)	(2, 2)	(1, 2)	(2, 0)	(0, 1)	(2, 1)	(0, 2)	(1, 0)
(1, 2)	(0, 0)	(1, 2)	(2, 1)	(2, 2)	(0, 1)	(1, 0)	(1, 1)	(2, 0)	(0, 2)
(2, 0)	(0, 0)	(2, 0)	(1, 0)	(0, 1)	(2, 1)	(1, 1)	(0, 2)	(2, 2)	(1, 2)
(2, 1)	(0, 0)	(2, 1)	(1, 2)	(1, 1)	(0, 2)	(2, 0)	(2, 2)	(1, 0)	(0, 1)
(2, 2)	(0, 0)	(2, 2)	(1, 1)	(2, 1)	(1, 0)	(0, 2)	(1, 2)	(0, 1)	(2, 0)

(9.81)

■

EXERCISE 9.15.14 Construct a field with eight elements, providing complete Cayley tables for addition and multiplication.

Index